

HARVARD UNIVERSITY
Graduate School of Arts and Sciences



DISSERTATION ACCEPTANCE CERTIFICATE

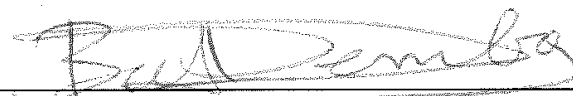
The undersigned, appointed by the
Harvard John A. Paulson School of Engineering and Applied Sciences
have examined a dissertation entitled:
“Algorithms and Array Design Criteria for Robust Imaging in
Interferometry”
presented by: Binoy George Kurien
candidate for the degree of Doctor of Philosophy and here by
certify that it is worthy of acceptance.

Signature 

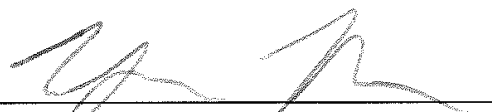
Typed name: Professor V. Tarokh

Signature 

Typed name: Professor T. Zickler

Signature 

Typed name: Professor D. Ba

Signature 

Typed name: Dr. Y. Rachlin

Date: April 20, 2016

Algorithms and Array Design Criteria for Robust Imaging in Interferometry

A dissertation presented

by

Binoy George Kurien

to

The Harvard John A. Paulson School of Engineering and Applied Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Engineering Sciences

Harvard University

Cambridge, Massachusetts

April 2016

© 2016 Binoy George Kurien

All rights reserved.

Dissertation Advisor:
Professor Vahid Tarokh

Author:
Binoy George Kurien

Algorithms and Array Design Criteria for Robust Imaging in Interferometry

Abstract

Optical interferometry is a technique which can be used to obtain high-resolution imagery of a distant target by interfering light collected by multiple telescopes. Image restoration from interferometric measurements poses a unique set of challenges. The first challenge is that the measurement set provides only a sparse-sampling of the object's Fourier Transform and hence image formation from these measurements is an inherently ill-posed inverse problem. Secondly, atmospheric turbulence causes severe distortion of the phase of the measured Fourier samples. We develop array design conditions for unique recovery of the true Fourier phase in the presence of this distortion, as well as a comprehensive algorithmic framework based on the notion of redundant-spaced-calibration (RSC), which together achieve reliable image reconstruction in spite of these challenges. Within this framework, we see that the classical interferometric observables known as the *bispectrum* and *closure phase* can limit sensitivity, and that generalized notions of these observables can improve both theoretical and empirical performance in Fourier phase estimation. Our framework leverages techniques from lattice theory to resolve integer phase ambiguities in the interferometric phase measurements, and from graph theory, to select a reliable set of generalized observables. As part of the performance assessment of our algorithm, we first show that both the theoretical and simulated Fourier-phase estimation accuracy of the algorithm approach an atmosphere-oracle Cramér-Rao Lower Bound at flux levels as low as 60 photons per interferometric fringe in the shot-noise-limited setting. Leveraging techniques from the field of *sparse recovery*, we then demonstrate reliable image reconstruction from the recovered Fourier estimates. Our results show that reconstructed image quality is retained

even in simulated stressing scenarios consisting of per-exposure flux levels on the order of 10 photons per interferometric fringe with less than 10 minutes of observation time. The end result is a comprehensive strategy to achieve reliable image reconstruction of dim objects with optical interferometry.

Contents

Abstract	iii
Acknowledgments	xii
Introduction	1
1 Fundamentals of Optical Interferometry	6
1.1 Chapter Overview	6
1.2 Interferometric Architectures	7
1.3 The Van-Cittert-Zernike Theorem	8
1.4 The Cramer-Rao Bound for the Complex Visibilities	14
1.5 Key Parameters of an Interferometer: Resolution, Field-of-View, and Under-sampling Ratio	17
1.6 Atmospheric and Instrumental Phase Noise	19
2 Leveraging compressed sensing techniques in optical interferometry	22
2.1 Citation to Previously-Published Work	22
2.2 Chapter Overview	22
2.3 Problem Statement and Approach	23
2.4 Description of Stage 1 of Algorithm	26
2.5 Description of Stage 2 of Algorithm	28
2.6 Algorithm Performance	30
2.6.1 Laboratory Validation	30
2.6.2 Simulation	30
3 Pattern design criteria for uniqueness in phase recovery	34
3.1 Citation to Work under Review	34
3.2 Chapter Overview	34
3.3 Preliminaries	35
3.3.1 Lattices	35
3.3.2 The Closest-Vector-Problem	36
3.4 Problem Statement and Related Work	37

3.5	Phase Wrapping Ambiguities in RSC Image Reconstruction	42
3.5.1	Identifying the Fundamental Phase Ambiguity	42
3.5.2	Quantifying the Effect of the Fundamental Ambiguity	49
3.5.3	Relation to closure-phase approaches	51
3.6	Implications of Wrap Ambiguities on Pattern Design	56
3.7	Wrap-invariance and Practical RSC Calibration	62
3.7.1	Practical Phase Approaches	62
3.7.2	Practical Phasor Approaches	66
3.8	Conclusions	69
4	Robust image reconstruction with redundant arrays and generalized closure phases	71
4.1	Chapter Overview	71
4.2	Problem Statement and Prior Work	72
4.3	Preliminaries	78
4.3.1	Fringe Noise Model for Pairwise Beam Combiner	78
4.3.2	N-Spectrum Covariance and SNR Models for the Pairwise Architecture	80
4.3.3	Covariance Matrix of the Generalized Closure Phases for the Pairwise Architecture	82
4.3.4	Covariance Matrix Approximations for the Fizeau Architecture	84
4.4	Fourier Phase Recovery using the N-Spectra	85
4.4.1	RSC Phase Recovery with Wrap-Invariant Closure Mappings	85
4.4.2	Selection of the N-Spectra	93
4.5	A Practical Algorithm for RSC Closure Imaging	96
4.6	Algorithm Performance	97
4.6.1	Sensitivity Limits	97
4.6.2	Simulation	99
4.6.3	Generalized Closures in Non-Linear Least Squares Approaches	108
4.7	Conclusion	113
5	Conclusions	114
	References	117
	Appendix A Appendix to Chapter 1	122
A.1	Sinusoidal Dependence of Field on Interferometer Focal Plane	122
	Appendix B Appendix to Chapter 3	124
B.1	Proof of Proposition 3.5.7	124
B.2	Proof of Lemma 3.7.2	125

Appendix C	Appendix to Chapter 4	128
C.1	Fizeau Variance Approximations	128
C.1.1	Variance Decomposition	128
C.1.2	Order-2 Terms	131
C.1.3	Higher Order Terms	133
C.2	Proofs of Proposition 4.4.2 and Corollary 4.4.3	134
C.2.1	Proof of Proposition 4.4.2	134
C.2.2	Proof of Corollary 4.4.3	134
C.3	Proof of Proposition 4.4.5	135
C.4	Minimum Cycle Basis Proofs	138

List of Tables

5.1	Summary of new uniqueness results presented in this thesis: the implications of wrap-invariant RSC patterns	116
-----	---	-----

List of Figures

1.1	The two popular beam combination schemes in optical interferometry . . .	8
1.2	The Fizeau Interferometer Concept	9
1.3	Path difference between two apertures	11
1.4	The relationship between aperture patterns and Fourier sampling	17
1.5	Representation of a scene as a Field-of-View comprised of resolution elements	18
1.6	Eliminating the effect of atmospheric turbulence with phase closure	21
2.1	Overview of Proposed Two-Stage Approach	26
2.2	Inverted Gaussian penalty function ($1 - \mathbf{h}_{\text{gaussian}}$); the dark region represents an area of near-zero penalty	28
2.3	Non-redundant Golay 20-aperture pattern (left) and corresponding UV- sampling (right)	31
2.4	Target chrome mask (left) and corresponding <i>truth</i> image at diffraction-limited resolution (right)	31
2.5	Image Reconstruction Results from Laboratory Validation	31
2.6	Image Reconstruction Results from Simulations	32
2.7	Image Reconstruction Results from Simulations	33
2.8	Algorithm Performance in Simulation	33
3.1	Fraction of redundant baselines required for critical redundancy vs. aperture count	38
3.2	Y-pattern Array Example	40
3.3	Reconstruction Results for Y-Pattern Example	41
3.4	Example: Dependent <i>Parallelogram</i> redundancy.	44
3.5	Example: 5-aperture RSC pattern. The six distinct baselines are shown. . . .	44
3.6	Illustration of the fundamental ambiguity of 2π -periodicity in RSC imaging. Distinct unwrappings $\hat{\beta}^*$ and β_0^* both produce solutions to Equation (3.5) in the noiseless case, as do their respective projections onto $K + L$, $\hat{\beta}_{K+L}^*$ and $\hat{\beta}_{0,K+L}^*$, in the noisy case.	49
3.7	Distinction between spanning tree baselines (thick, solid) and loop entry baselines (thin, dotted)	53

3.8	Bootstrapping phase of a low-SNR baseline (green) with subset (blue) of high-SNR baselines from spanning tree baselines (black)	55
3.9	Example: Reducing an aperture pattern and associated matrix to identify Persistent Loop(s)	58
3.10	Persistent Loop set at center of pattern in Figure 3.2	60
3.11	Amended Pattern	61
3.12	Reconstruction Results for Amended Pattern	61
3.13	Reconstruction Results for Phasor Approach	69
4.1	Generalizing the phase closure concept . SNRs given assume each aperture contributes $\hat{n} = 2e3$ photoelectrons to each fringe in pairwise combination. .	76
4.2	One possible cycle-basis for a simple interferometric graph. Note the fourth triangle (i.e. baseline set $\{2, 4, 6\}$) can be expressed as a linear combination of the cycles shown.	87
4.3	A shortest path tree	95
4.4	RSC aperture pattern used in simulation	101
4.5	UV-sampling for RSC pattern	101
4.6	Golay non-redundant beam-combiner pattern	102
4.7	Truth image for simulation: the CALIPSO satellite	102
4.8	Truth image at the resolution of the interferometric pattern	103
4.9	Pairwise Phase Recovery Results for Flux $n = 2e3$ pe/ap/frame, $5e4$ frames	105
4.10	Fizeau Phase Recovery Results for Flux $n = 2e3$ pe/ap/frame, $5e4$ frames .	105
4.11	Fizeau Phase Recovery Results for Flux $n = 2e3$ pe/ap/frame, $1e3$ frames .	106
4.12	Fizeau Phase Recovery Results for Flux $n = 5e2$ pe/ap/frame, $5e4$ frames .	106
4.13	Fizeau Image Reconstruction Results. The reconstructions in the left column used traditional, three-baseline observables, whereas those in the right column used generalized closures selected according to Algorithm 3. (top row) $n = 2e3$ photoelectrons/aperture/frame (pe/ap/frame), $5e4$ frames, (middle row) $n = 2e3$ pe/ap/frame, $2e3$ frames, (bottom row) $n = 5e2$ pe/ap/frame, $5e4$ frames.	107
4.14	(top row) Reconstructed images and convergence times for NLS algorithm with traditional-closure basis (left), generalized-closures from minimum cycle basis (middle), and generalized-closures from minimum-variance spanning tree (right), and (bottom row) Reconstructed images and elapsed running times with all $\binom{N_{ap}}{3}$ traditional closures at iteration 20 (left), iteration 40 (middle), and iteration 80 (right)	110
4.15	Raw (left) and median-filtered (right) Fourier phase error using the NLS approach.	111

4.16 The minimum-variance spanning tree	112
A.1 Beam-Combination Geometry	123
C.1 Illustrating Lemma C.4.1 (a) and Lemma C.4.2 (b)	139

Acknowledgments

It is with great pleasure that I thank my thesis advisor Prof. Vahid Tarokh. Vahid offered me complete freedom in choosing a path for my dissertation research, and then conferred the wisdom necessary to maintain steady progress on this path. Along the way, his sharp insight often revealed the mathematical or algorithmic source of an obstacle to my progress, thereby helping me to proceed efficiently towards the results described in this document. As a matter of equal importance from my own perspective, I feel I have become a better engineer and mathematician over the past four years, and this improvement is in large part due to the sincerity and sharpness of the feedback that Vahid provided. For this I am also very grateful.

It also gives me great pleasure to thank all of my colleagues at MIT Lincoln Laboratory (MIT/LL). Dr. Jonathan Ashcom, Dr. Yaron Rachlin, and Vinay Shah warmly welcomed me onto their Interferometry team three years ago. Through collaboration with them, I was able to learn the fundamentals of optical interferometry and help develop new algorithms in a friendly and intellectually-stimulating environment. In addition to serving admirably as the leader of the team, Jon Ashcom went out of his way on numerous occasions to aid my research progress on a myriad of fronts ranging from intellectual to logistic and administrative. I am grateful for both Jon's encouragement and abundant technical insights throughout my time on the project. Discussions with Yaron were invaluable in defining the directions for my theoretical research at certain critical junctures in the PhD process. I have always appreciated Yaron's ability to clearly identify theoretical work that is worth pursuing in earnest, and am fortunate that he availed this talent to help me further my research. Working with Vinay and the rest of the Interferometry team in reconstructing imagery from data collected by the MIT/LL Interferometer has been a fun and rewarding experience. I am also grateful for numerous helpful conversations with Vinay. He has patiently listened to my algorithmic ideas, including many that were, in hindsight, either crazy or flawed (and sometimes both). To these admittedly-quixotic ideas, he has replied not with the complete dismissal or derision they may have deserved, but rather with constructive and

highly-insightful suggestions on how to project the ideas onto a more "realistic subspace". I would also like to thank MIT/LL Group Leader Dr. Sumanth Kaushik for granting me the opportunity to collaborate on the Interferometry project, and for his continued encouragement of my research.

I am extremely grateful to have been granted fellowship support for my PhD through the MIT/LL Lincoln Scholars Program. I thank current and former members of the Lincoln Scholars committee, and especially Dr. Mark Weber, Dr. James Ward, and Mr. Ken Estabrook, for their continued support throughout this Program. I am grateful to fellow members of Group 64 at MIT/LL, and especially Mrs. Linda Weeks, Ms. Lori Jeromin, and current and former Group Leaders Dr. Thomas Macdonald and Ms. Leslie Alger, respectively, for supporting my growth as an engineer and my admission into the Program.

I would like to thank the faculty and staff at Harvard SEAS for an enriching four years of study. First, I would like to acknowledge Prof. Todd Zickler and Prof. Demba Ba for serving on my PhD Research Committee. Secondly, it was a privilege to take courses in Harvard SEAS and to learn from the esteemed Harvard faculty. In particular, I would like to thank Prof. Yue Lu. I was very fortunate to be enrolled in the Statistical Inference course taught by Prof. Lu. This class provided me with a solid foundation in Estimation and Detection Theory which was critical for my subsequent research. Prof. Lu's dedication to maximizing the learning experience for his students was evident in the careful, structured presentation style in his lectures, the clarity of which seemed to be near-optimal with remarkable consistency! I am also thankful for the support and assistance of the staff at Harvard SEAS, and especially to Ms. Kathy Masse and Ms. Lisa Frazier-Zezze. Both Kathy and Lisa were always so pleasant and helpful to me, and I especially appreciated this during otherwise-stressful occasions including my qualifying examination.

The company of my family and friends kept me sane and happy in the course of the PhD process. I am truly grateful for the support of all of my friends and family members. I will mention a few with whom the past four years will be indelibly linked in my memory. My Aunt Betsy and my Uncle David visited many times from Connecticut, bearing delicious

Bacon-Lettuce-and-Tomato (BLT) sandwiches and their great company. I am grateful to them for these peaceful interludes. I am also grateful to my dear friends Dheeraj and Shruti, with whom my wife Neha and I have enjoyed many weekend outings which have been relaxing, memorable, and often quite comical. My mom Jean has been a steady source of encouragement and support throughout my life in times of success as well as in times of failure. My dad Tom has always encouraged diligence in my education and in my career, and has instilled in me the importance of being systematic in engineering, writing, and presenting. My brother Joe has provided welcome perspective and many good laughs over the last few years.

Finally, and most importantly, I would like to thank my dear wife Neha for her loving willingness to accompany and support me in this journey. Her companionship and natural cheerfulness supplied quite a formidable force of happiness to counter the stresses of the last four years. Neha, I am so lucky to have you as my best friend.

ACKNOWLEDGEMENT: This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering.

This thesis is dedicated to my parents, Jean and Tom Kurien.

Introduction

The use of optical interferometry as a multi-aperture imaging approach is attracting increasing interest in the astronomical and remote-sensing communities. The appeal of this technique is primarily due to the high resolution it affords relative to single-aperture imaging. Namely, the diffraction-limited angular resolution of a single aperture is proportional to $\frac{\lambda}{D}$, where λ is the wavelength of the light, and D is the diameter of the aperture. On the other hand, the achievable angular resolution of an interferometric array of apertures is instead proportional to $\frac{\lambda}{B_{max}}$, where B_{max} is the maximum spatial separation of any two apertures in the array. Therefore with interferometry one can achieve the same high resolution offered by an extremely large (and often prohibitively-costly) telescope by interfering light from several telescopes of practical size distributed over a large area.

Image reconstruction in interferometry is akin to the more general problem of imaging with sparse samples in the Fourier domain, which is encountered in many other fields from medical imaging to radio astronomy. Namely, the fundamental problem is to regularize the ill-posed reconstruction of N_{resEl} resolution elements in the astronomical scene from $m \ll N_{resEl}$ Fourier samples derived from the array's interference patterns. Traditionally approaches based on the so-called *CLEAN* algorithm due to Högbom (1974), which implicitly fit the measurements to an image model consisting of a sparse collection of point-like sources, have been used to solve this problem. Given the success of this relatively-simple algorithm, it is hardly surprising that recently-developed techniques leveraging sophisticated sparse image models (Wiaux *et al.*, 2009) (Kurien *et al.*, 2014) have shown promise. In fact such algorithms belong to a burgeoning family of cross-disciplinary techniques, which are based

on seminal work within the last decade (Donoho, 2006) (Candès *et al.*, 2006), in the allied fields of *sparse recovery* and *compressed sensing*. These algorithms generally take advantage of the fact that while undersampled Fourier measurement sets do not uniquely define an arbitrary underlying image, they can uniquely specify a special set of images of practical interest: those that are *sparse* in some subspace.

If Fourier samples were directly available, we could directly apply standard techniques for image recovery from a sparse measurement set in the frequency domain. However, estimation of the Fourier samples is itself a complicated inference problem. Interferometric systems interfere signals from multiple apertures to produce superpositions known as *fringes* which encode samples of the Fourier Transform of the object under observation. In our introductory Chapter 1, we present a mathematical model for fringe observation as well as sensitivity limits for estimation of the fringe parameters. Estimates of the scene’s brightness-normalized Fourier samples, which are known as *complex visibilities*, can then be derived from these parameters. A fundamental challenge in interferometry is the distortion of the phase of the measured complex visibilities due to natural variation in the effective path lengths to the target observed by each aperture. In practice, the most stressing source of this variation is turbulence in the Earth’s atmosphere. This turbulence alters the mean phase profile at each aperture in the array by a non-uniform and time-varying amount (i.e. the so-called *optical piston*), which in turn causes rapid shifting of the fringes on the focal plane. If uncorrected, the resulting phase noise in the Fourier estimates then causes severe distortion in the reconstructed image.

Techniques for elimination of this phase noise fall into two categories according to the photon flux level. At high photon-flux levels, the fringes are bright and the complex visibility estimates have high Signal-to-Noise ratio (SNR). In this case, it is possible to directly decouple the contributions of the true Fourier phase from the confounding phase contribution arising from the unknown, atmosphere-induced piston variation. While the problem is ill-posed in general, this decoupling can be performed by alternating estimation of the Fourier and image pixels while enforcing a-priori constraints in the image domain,

such as source-sparsity (Pearson and Readhead, 1984). At low flux levels, on the other hand, the individual fringe exposures (or *frames*) are typically too weak for such techniques to be reliable. Moreover, the presence of random phase variation across frames due to the atmosphere means that these fringe measurements cannot be directly integrated. Instead atmosphere-invariant derivatives ¹ of the fringe measurements such as the bispectrum and the power spectrum are integrated and the complex visibilities must then be inferred from these integrated observables.

This thesis in large part reflects the confluence of two emerging trends in interferometry and in inverse imaging problems taken collectively: in the former, there is a need for approaches for imaging complex, extended objects with a sparse measurement set and limited a-priori knowledge, and in the latter, sparse recovery and compressed sensing (CS) techniques continue to show great promise in imaging a vast range of natural images from similarly-sparse measurement sets. While standard CS techniques are tailored to linear-inverse problems and do not apply to atmosphere-invariant interferometric observables, we develop and validate an algorithmic interface in Chapter 2 permitting these techniques to be used successfully. We show empirically that for compact scenes the effect of the atmosphere can be eliminated for a wide range of fluxes. The limiting restriction of our interface to the imaging of compact scenes is a direct consequence of the fundamental ill-posed nature of Fourier phase recovery from atmosphere-perturbed fringe measurements. Namely, since the mapping from Fourier phase to the set of measured phases is rendered non-injective ² by the atmosphere in general, we must introduce constraints upon the object to uniquely determine the Fourier phase. Such regularization techniques and the object constraints they impose can be obviated by introducing redundancy into the inter-aperture spacings in the interferometric array. This well-posed framework for Fourier phase determination forms the basis of Chapters 3 and 4.

¹By *derivative*, we mean an observable derived from the fringe measurements, and not the mathematical derivative.

²By *non-injective*, we simply mean that a given measurement set can be produced by different sets of Fourier phases so that the former does not uniquely determine the latter.

The use of redundant aperture spacing to recover the Fourier phase in the presence of piston variation in a well-posed manner is known as *redundant spacing calibration (RSC)*. Of the myriad RSC techniques that have been developed for both radio and optical interferometry, some operate on the measured complex visibilities (which we call *Phasor* approaches), and others on the phase component of these complex visibilities (which we call *Phase* approaches). Even with redundancy, a fundamental ambiguity exists in RSC-based phase recovery which is rooted in the 2π -periodicity of the interferometric phase. As a direct result of this ambiguity, the solution for the complex visibilities given the measured visibilities is non-unique in both Phase and Phasor approaches. However, we will show in Chapter 3 that for certain patterns, which we denote *wrap-invariant* patterns, this fundamental ambiguity can be rendered benign. Namely, for such patterns, phase-unwrapping techniques such as those based on the closest-vector-problem (CVP) formulation due to Lannes and Anterrieu (1999) can successfully recover the true Fourier phase (modulo 2π). We show that this wrap-invariance property is conferred upon arrays whose interferometric graph satisfies a certain cycle-free condition. This condition is, to the best of our knowledge, the first sufficient condition on an interferometric aperture pattern for unique recovery of the Fourier phase. For cases in which this condition is not satisfied, we provide a simple algorithm for identifying those graph cycles which prevent its satisfaction. For illustrative purposes, we apply this algorithm to diagnose a member of a aperture-pattern family popular in the literature which is not wrap-invariant, and modify it so that it achieves wrap-invariance.

Having established conditions for unique Fourier phase recovery, we then turn our attention in Chapter 4 to the issue of well-posed and practical image formation. In particular, we consider low-flux scenarios in which atmosphere-invariant observables must be integrated over many frames for reliable phase estimation. In hopes of improving sensitivity, we generalize the classical atmosphere-invariant phase observables (i.e. the bispectrum and closure phase) to higher-order observables, which we denote as the *n-spectrum* and *generalized closure phase* respectively. We extend the uniqueness results in Chapter 3 to our new observables, and develop a novel comprehensive algorithm for image reconstruction

from these observables. Here we leverage the notion of *minimum cycle basis* from graph theory. A standard *sparse-recovery* technique known as *total-variation minimization* is used to perform the final image reconstructions. We present numerical and visual evidence indicating that our generalized observables can yield better estimation performance relative to their classical counterparts. Moreover, we show that our algorithm's performance approaches the Cramer-Rao Lower Bound for this estimation problem in an example scenario of imaging a dim object.

In summary, great progress has been made recently in solving standard, linear inverse problems in imaging as part of the development of *compressed sensing* and *sparse recovery* methodology. Simultaneously, many algorithmic techniques have been developed for solving the specialized ill-posed inverse problem in optical interferometry, which is greatly complicated by the atmosphere. Characterizing the penalty in achievable performance associated with this increased complexity has remained an open problem. Hence in this thesis we ask the question: "To what extent can we make interferometric imaging immune to the effects of the atmosphere in both reliability and sensitivity?". We then probe the limits of this immunity using a novel algorithmic framework which guarantees unique phase recovery through the notion of wrap-invariance, and generalizes the notion of atmosphere-invariant observable in order to approach theoretical sensitivity limits at low flux levels. In these two pursuits, our framework leverages theory and algorithms from lattice theory and graph theory, respectively. In the end we provide strong evidence that even in scenarios thought to be stressing from an SNR perspective, we can achieve near-immunity to the effects of the atmosphere.

Chapter 1

Fundamentals of Optical Interferometry

1.1 Chapter Overview

In this chapter, we introduce the physics-based principles of optical interferometry, thereby providing a foundation for the main results of this thesis. As the cornerstone of these principles, the Van-Cittert-Zernike Theorem establishes the interferometer as a *Fourier-imaging device*, i.e. one whose output can be easily-conceptualized when examined in the Fourier domain. In particular, we show that interferometric arrays characterize the object under observation via measurement of its brightness-normalized Fourier components, which are the so-called *complex visibilities*. We then provide a derivation of the sensitivity limits for measurement of these complex visibilities in the shot-noise-limited case. This derivation is based on the well-known Cramer-Rao Lower Bound (CRLB) in estimation theory, and corroborates previous results due to Zmuidzinas (2003). This CRLB will prove useful as a performance benchmark for the algorithms developed in this thesis. We then proceed to introduce the role of atmospheric turbulence and other sources of interferometric phase noise in interferometry. Finally we define a few important parameters of an interferometer: *resolution element*, *Field-of-View*, and *undersampling ratio*.

1.2 Interferometric Architectures

The main appeal of interferometry over other astronomical imaging techniques is that it affords the same high-resolution capability offered by an extremely large (and often prohibitively-costly) telescope by interfering light from several telescopes of practical size. Optical interferometers enable the imaging of a distant scene by providing a sampling of the scene's 2D Fourier Transform. Several excellent surveys provide the theoretical basis of interferometry as well as the practical issues involved in building and operating an interferometer, including those by Labeyrie *et al.* (2006), Glindemann (2011), and Buscher (2015). Each pair of telescopes in an interferometric array measures a single angular spatial frequency of $\frac{2\pi\mathbf{b}}{\lambda}$ radians, where \mathbf{b} is the vector difference of the telescope positions, which is known as a *baseline*. For an array of N_{ap} apertures, the data set then consists of all $\binom{N_{ap}}{2}$ such measurements.

There are two popular beam combination architectures in use in optical interferometry: the pairwise combination scheme, and the Fizeau combination scheme. The two schemes are illustrated in Figure 1.1. Suppose we have a telescope array consisting of N_{ap} apertures, each of which collects n photons from a distant source. In the Fizeau scheme, light from all telescopes is interfered on a single focal plane, forming an interference pattern known as a *fringe* for each telescope pair. Each fringe encodes a sample of the scene's 2D Fourier Transform. In the pairwise scheme, light from each aperture is split N_{ap} ways and combined with that of each other aperture to form a single fringe pattern on separate focal planes. Hence each of the $\binom{N_{ap}}{2}$ focal planes receives $\frac{2n}{N_{ap}-1}$ photons, which is typically a small fraction of the total photons incident upon the array.

Ideally an interferometer¹ provides a perfect encoding of the scene's sampled Fourier Transform. In practice, however, we never observe pristine interference fringes in either architecture. The quality of the *fringes* is degraded by statistical fluctuations in the arrival times of photons on the focal plane, which is a phenomenon known as *shot noise*. We will

¹We will use the terms *interferometer* and *interferometric array* interchangeably in this thesis

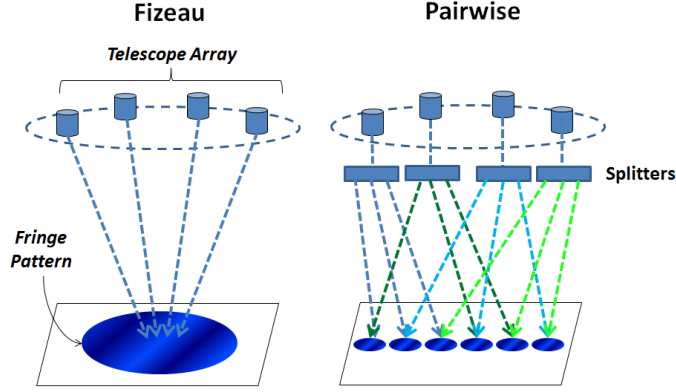


Figure 1.1: The two popular beam combination schemes in optical interferometry

quantify the impact of shot noise further in the course of the thesis. For now, we note that the Fizeau architecture uses all light to form each fringe while incurring shot noise from all $N_{ap}n$ photons. On the other hand, the pairwise architecture uses a fraction of the light to form each fringe, which is then by corrupted by the shot noise due to the photons collected by two apertures. A fundamental result due to Zmuidzinas (2003) quantified this tradeoff and established the superior overall sensitivity of the Fizeau (or *all-in-one*) scheme with respect to the pairwise scheme. Because the Fizeau architecture is also typically simpler to implement, it is often preferred in practice. The results in this thesis are therefore geared toward the Fizeau architecture, although we will leverage the pairwise architecture in Chapter 4 as a conceptual springboard for subsequent analysis.

1.3 The Van-Cittert-Zernike Theorem

To derive the Theorem, let us model a source in the sky as a superposition of differential emitting elements of size $\Delta\Omega$. Consider then the electric field at the j -th aperture of an interferometric array which is a distance d from one such differential patch. This can be written as:

$$E_j(\Omega) = \frac{a(\Omega)\Delta\Omega}{R_j} e^{i\omega(t - \frac{R_j}{c})} \quad (1.1)$$

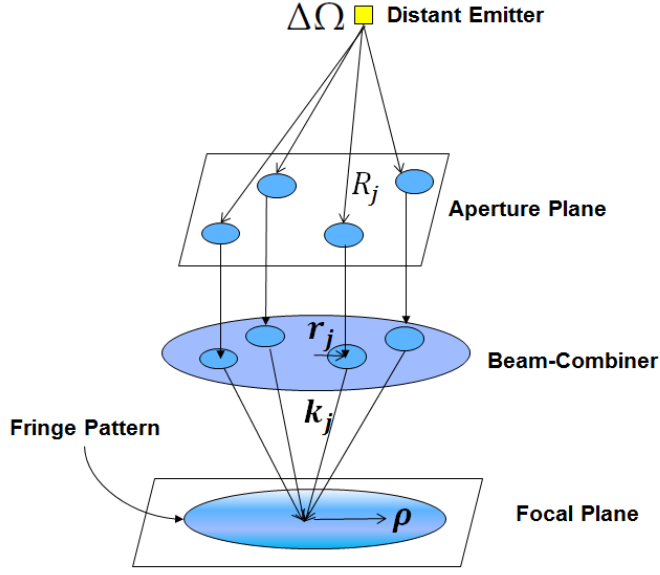


Figure 1.2: *The Fizeau Interferometer Concept*

where $a(\Omega)$ is the amplitude of the field at the source, R_j is the distance from the object to the j -th aperture, ω is the angular frequency of the wave, and c is the speed of light. For the purposes of this analysis, we neglect the amplitude variation amongst elements of the array, so that the amplitude at all apertures is a constant $\frac{a(\Omega)\Delta\Omega}{d}$.

The fields from all of the point sources in the interferometer's field-of-view are superimposed on the system's focal plane. For the *Fizeau* interferometer, this superposition is typically performed by a beam-combiner, which focuses the light collected by each aperture to a common point with propagation direction specified by an aperture-specific wave-vector \mathbf{k}_j (see Figure 1.2).

Using the definitions above, we can compute the field contribution of each aperture at an arbitrary point on the focal plane $\boldsymbol{\rho}$. This contribution will be a function of the respective wave-vectors and vector distances from $\boldsymbol{\rho}$. Namely, we have:

$$E_j(\boldsymbol{\rho}) = \frac{a(\Omega)\Delta\Omega}{d} \cdot e^{i(\mathbf{k}_j \cdot \mathbf{x}_j(\boldsymbol{\rho}))} e^{i\omega(t - \frac{R_j}{c})} \quad (1.2)$$

where \mathbf{k}_j is the wavevector of the outgoing wave from aperture j at the beam-combiner, and $\mathbf{x}_j(\boldsymbol{\rho})$ is the vector position to $\boldsymbol{\rho}$ relative the j -th aperture position at the beam-combiner.

As we show in Appendix A.1, we can rewrite this expression in terms of the lateral aperture position \mathbf{r}_j at the beam combiner as:

$$E_j(\boldsymbol{\rho}) = \frac{a(\Omega)\Delta\Omega}{d} \cdot e^{i(\mathbf{r}_j \cdot \boldsymbol{\rho} + \phi_j)} e^{i\omega(t - \frac{R_j}{c})} \quad (1.3)$$

where ϕ_j is an aperture-dependent phase which is independent of $\boldsymbol{\rho}$.

Now consider the field superposition at the pixel at vector coordinate $\boldsymbol{\rho}$ on the focal plane (see Figure 1.2), which can be written as:

$$E(\Omega) = \frac{a(\Omega)\Delta\Omega}{d} \cdot \left(\sum_{j=1}^{N_{ap}} e^{i(\mathbf{r}_j \cdot \boldsymbol{\rho} + \phi_j)} e^{i\omega(t - \frac{R_j}{c})} \right) \quad (1.4)$$

where we have neglected any attenuation incurred by the field en-route to the focal plane for simplicity. To obtain the field for the complete source, we now integrate over all differential patches (which we call *point sources*) in the source to obtain:

$$E_{total} = \int d\Omega \cdot \frac{a(\Omega)}{d} \left(\sum_{j=1}^{N_{ap}} e^{i(\mathbf{r}_j \cdot \boldsymbol{\rho} + \phi_j)} e^{i\omega(t - \frac{R_j}{c})} \right) \quad (1.5)$$

The detector at pixel $\boldsymbol{\rho}$ measures the intensity Q of the superposition. Namely,

$$Q(\boldsymbol{\rho}) = |E_{total}|^2 = E_{total} E_{total}^* \quad (1.6)$$

or,

$$Q(\boldsymbol{\rho}) = \frac{1}{d^2} \int d\Omega_1 \cdot a(\Omega_1) \left(\sum_{j=1}^{N_{ap}} e^{i\omega(t - \frac{R_j}{c})} e^{i(\mathbf{r}_j \cdot \boldsymbol{\rho} + \phi_j)} \right) \int d\Omega_2 \cdot a^*(\Omega_2) \left(\sum_{k=1}^{N_{ap}} e^{-i\omega(t - \frac{R_k}{c})} e^{-i(\mathbf{r}_k \cdot \boldsymbol{\rho} + \phi_k)} \right)$$

The final measurement is a photon count which is proportional to the time-averaged intensity. Hence, after re-arranging and time-averaging, we obtain:

$$\langle Q(\boldsymbol{\rho}) \rangle = \frac{1}{d^2} \iint \langle a(\Omega_1) a^*(\Omega_2) \rangle \left(\sum_{j=1}^{N_{ap}} e^{i\omega(t - \frac{R_j}{c})} e^{i(\mathbf{r}_j \cdot \boldsymbol{\rho} + \phi_j)} \right) \left(\sum_{k=1}^{N_{ap}} e^{-i\omega(t - \frac{R_k}{c})} e^{-i(\mathbf{r}_k \cdot \boldsymbol{\rho} + \phi_k)} \right) d\Omega_1 d\Omega_2 \quad (1.7)$$

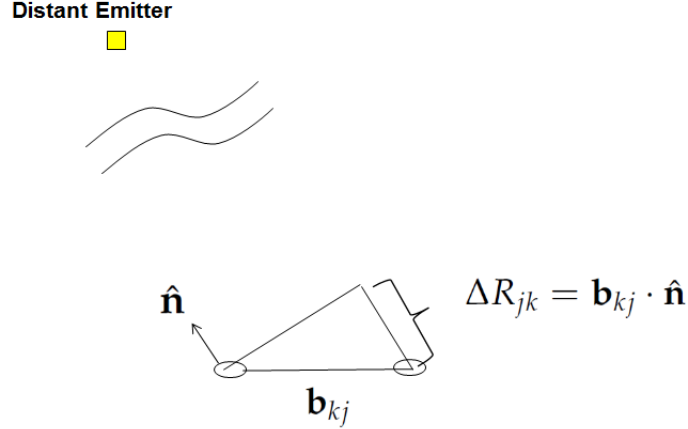


Figure 1.3: Path difference between two apertures

We can simplify this expression greatly by noting that, with a few exceptions, astronomical sources are *spatially-incoherent*, which means that the complex amplitudes of radiated waves from different points are uncorrelated. Namely,

$$\langle a(\Omega_1)a^*(\Omega_2) \rangle = I(\Omega)\delta(\Omega_1 - \Omega_2) \quad (1.8)$$

where $\delta(x)$ denotes the Dirac delta function.

The result is that the double integral above collapses to a single integral in Ω . We can now write the sum of all pairwise products in Equation (1.7) as the following single summation over all $\binom{N_{ap}}{2}$ aperture pairs:

$$\begin{aligned} \langle Q(\boldsymbol{\rho}) \rangle &= \frac{1}{d^2} \int I(\Omega) N_{ap} d\Omega + \\ &\frac{1}{d^2} \sum_{(j,k): k>j}^{N_{ap}} (e^{i(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_{jk} + \phi_{jk})} \int I(\Omega) e^{i(\frac{\omega \Delta R_{jk}}{c})} d\Omega + e^{-i(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_{jk} + \phi_{jk})} \int I(\Omega) e^{-i(\frac{\omega \Delta R_{jk}}{c})} d\Omega) \end{aligned} \quad (1.9)$$

where the first term represents the sum of the products of each aperture's field with its conjugate (i.e. terms of the form $E_j E_j^*$).

To understand the meaning of Equation (1.9), it is useful to approximate ΔR_{jk} using the geometry illustrated in Figure 1.3. Let us define the unit vector originating at aperture 1 and pointing in the direction of the emitting source as $\hat{\mathbf{n}}$. For an arbitrary source location,

this unit vector would of course differ for each aperture. In most astronomical scenarios, however, the inter-aperture distance is extremely small compared with the distance between the source and the aperture array (Buscher, 2015). In such cases it is reasonable to assume that the direction vector $\hat{\mathbf{n}}$ is approximately constant across the array. As shown in Figure 1.3, we have:

$$\Delta R_{jk} = \mathbf{b}_{jk} \cdot \hat{\mathbf{n}} \quad (1.10)$$

where \mathbf{b}_{jk} is the vector difference position of the k -th and j -th apertures.

Moreover, we can write the differential patch $d\Omega$ in terms of components of this vector \mathbf{n} as:

$$d\Omega = (d^2)dn_xdn_y \quad (1.11)$$

Substituting Equations (1.10) and (1.11) into Equation (1.9) and noting that $\frac{\omega}{c} = \frac{2\pi}{\lambda}$, we obtain:

$$\begin{aligned} \langle Q(\boldsymbol{\rho}) \rangle &= \int I(\Omega) N_{ap} dn_x dn_y + \\ &\sum_{(j,k): k > j}^{N_{ap}} (e^{i(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_{jk} + \phi_{jk})} \int I(\Omega) e^{i(\frac{2\pi \mathbf{b}_{jk}}{\lambda} \cdot \hat{\mathbf{n}})} dn_x dn_y + e^{-i(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_{jk} + \phi_{jk})} \int I(\Omega) e^{-i(\frac{2\pi \mathbf{b}_{jk}}{\lambda} \cdot \hat{\mathbf{n}})} dn_x dn_y) \end{aligned} \quad (1.12)$$

Note that the two integrals are $F^*(\frac{\mathbf{b}_{jk}}{\lambda})$ and $F(\frac{\mathbf{b}_{jk}}{\lambda})$, respectively, where $F(\frac{\mathbf{b}_{jk}}{\lambda})$ is the 2D Fourier Transform of the source evaluated at spatial frequency $\frac{\mathbf{b}_{jk}}{\lambda}$. For simplicity of notation, let us index all of the aperture pairs (j, k) with a single index h . Defining $\phi_h := \phi_{jk}$, and $F_h := F(\frac{\mathbf{b}_h}{\lambda})$ and substituting, we have:

$$\langle Q(\boldsymbol{\rho}) \rangle = \int I(\Omega) N_{ap} dn_x dn_y + \sum_{h=1}^{(N_{ap})} F_h e^{i(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_h + \phi_h)} + F_h^* e^{-i(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_h + \phi_h)} \quad (1.13)$$

$$\langle Q(\boldsymbol{\rho}) \rangle = F_0 N_{ap} + \sum_{h=1}^{(N_{ap})} 2|F_h| \cos(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_h + \angle F_h + \phi_h) \quad (1.14)$$

From this last equation, we see that focal plane consists of a superposition of 2D sinusoids

(or *fringes*), each of which encodes a distinct Fourier component of the source. This result is the Van-Cittert-Zernike Theorem which was first derived by Pieter van Cittert in 1934 (van Cittert, 1934) and then proved in simpler fashion by Frits Zernike in 1938 (Zernike, 1938).

For subsequent analysis in this thesis, it will be convenient to parameterize the fringes using the following Definitions:

Definition 1.3.1. The *complex fringe phasor* \mathbf{z} at a given spatial frequency is given by: $\mathbf{z} := |F_b|e^{j\angle F_b}$

Note that the magnitude of the fringe phasor will be proportional to the brightness of the object. It is also useful to have a quantity describing the strength of a given Fourier component relative to the overall brightness of the object. This relative measure is provided by the *complex visibility*:

Definition 1.3.2. The *complex visibility* \mathbf{v} of an object at a given spatial frequency is given by the ratio $\mathbf{v} = \frac{|F_h|e^{j\angle F_h}}{F_0}$

Definition 1.3.3. The *visibility* γ of an object at a given spatial frequency is the modulus of the complex visibility, i.e. $\gamma := \frac{|F_h|}{F_0}$

Definition 1.3.4. The *Fourier phase* θ of an object at a given spatial frequency is the argument (i.e. phase) of the complex visibility, i.e. $\theta := \angle F_h$.

With Definition 1.3.3, we can write Equation (1.15) as:

$$\langle Q(\boldsymbol{\rho}) \rangle = F_0 N_{ap} + \sum_{h=1}^{\binom{N_{ap}}{2}} 2\gamma_h F_0 \cos(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_h + \angle F_h + \phi_h) \quad (1.15)$$

The expected photon counts observed at the detectors are directly proportional to the intensity present. Let the proportionality constant relating intensity to photon counts be given by κ , so that a vectorized representation of these expected photon counts is given by:

$$\langle \mathbf{y}(\boldsymbol{\rho}) \rangle = \kappa F_0 N_{ap} + \kappa \sum_{h=1}^{\binom{N_{ap}}{2}} 2\gamma_h F_0 \cos(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_h + \angle F_h + \phi_h) \quad (1.16)$$

Moreover we assume that light is conserved throughout the fringe generation process, and hence the photon counts must sum to the total number T_p of photons incident upon the array, i.e. $T_p = nN_{ap}$, i.e.

$$T_p = \sum_{\rho} \langle \mathbf{y}(\rho) \rangle = \sum_{\rho} \left(\kappa F_0 N_{ap} + \kappa \sum_{h=1}^{\binom{N_{ap}}{2}} 2\gamma_h F_0 \cos(\rho \cdot \Delta \mathbf{r}_h + \angle F_h + \phi_h) \right) = nN_{ap} \quad (1.17)$$

In this thesis, we will consider the scenario in which the periods of the fringe sinusoids are all integer multiples of pixels on the focal plane. This is one of the so-called *DFT conditions* commonly employed in interferometry to avoid fringe estimation bias due to spectral leakage (Gordon, J. A. and Buscher, D. F., 2012). In this case the sinusoidal term vanishes in the sum in Equation (1.17) leaving only the constant (first) term, and if we define N_{fp} as the number of pixels on one side of a square focal plane, we have:

$$\sum_{\rho} \kappa F_0 N_{ap} = N_{fp}^2 \kappa F_0 N_{ap} = nN_{ap} \quad (1.18)$$

which in turn implies that $\kappa = \frac{n}{F_0 N_{fp}^2}$, and hence by substitution into Equation (1.16) yields:

$$\langle \mathbf{y}(\rho) \rangle = \frac{nN_{ap}}{N_{fp}^2} + \frac{2n}{N_{fp}^2} \sum_{h=1}^{\binom{N_{ap}}{2}} \gamma_h \cos(\rho \cdot \Delta \mathbf{r}_h + \theta_h + \phi_h) \quad (1.19)$$

1.4 The Cramer-Rao Bound for the Complex Visibilities

There are two principal sources of noise that affect this measurement. We have already alluded to the *shot noise* arising from the fact that photon arrivals are not uniform in time even if incident intensity is constant. The resulting uncertainty in photon counts for a given exposure time is well-modeled as a Poisson distribution with mean λT , where λ is proportional to the intensity and T is the exposure time. The second source of noise is *read noise*, which is thermal noise in the detection circuitry. Read noise is well-modeled by the Gaussian distribution. Since shot noise is proportional to intensity whereas read noise is

not, the former dominates the latter in high light-level scenarios. And with recent advances in electronics, the light-level regime in which read noise is important continues to diminish. We will therefore focus on detection limits for the shot-noise-limited case.

Recall Equation (1.19) provides an expression for the time-averaged intensity. To obtain lower bounds on sensitivity, let us assume that the phase differences $\Delta\theta_h$ are known and therefore calibrated. We can write the observation model in matrix form as:

$$\langle \mathbf{y}(\boldsymbol{\rho}) \rangle = \frac{nN_{ap}}{N_{fp}^2} + \frac{2n}{N_{fp}^2} \dots$$

$$\times \begin{bmatrix} \cos(\rho(1) \cdot \Delta \mathbf{k}_1) & -\sin(\rho(1) \cdot \Delta \mathbf{k}_1) & \cos(\rho(1) \cdot \Delta \mathbf{k}_2) & -\sin(\rho(1) \cdot \Delta \mathbf{k}_2) & \dots \\ \cos(\rho(2) \cdot \Delta \mathbf{k}_1) & -\sin(\rho(2) \cdot \Delta \mathbf{k}_1) & \cos(\rho(2) \cdot \Delta \mathbf{k}_2) & -\sin(\rho(2) \cdot \Delta \mathbf{k}_2) & \dots \\ \dots & & & & \end{bmatrix} \begin{bmatrix} \text{Re}[v_1] \\ \text{Im}[v_1] \\ \text{Re}[v_2] \\ \text{Im}[v_2] \\ \dots \end{bmatrix}$$

Suppose that the interferometer's output is a beam with an extent of N_{fp} pixels on the focal plane. Letting \mathbf{A} be the N_{fp}^2 -by- $2\binom{N_{ap}}{2}$ matrix in the equation above (the so-called *visibility-to-pixel matrix* (V2PM)), we can write this matrix equation compactly as:

$$\mathbf{y} = \frac{nN_{ap}}{N_{fp}^2} \mathbb{1}_{N_{fp}}^T + \frac{2n}{N_{fp}^2} \mathbf{A} \hat{\mathbf{v}} \quad (1.20)$$

where $\mathbb{1}_{N_{fp}}^T$ denote the constant *ones* vector of length N_{fp} , $\hat{\mathbf{v}}$ is the vector containing the quadrature components of the complex visibilities. The actual pixel counts can be modeled as i.i.d. Poisson random variables with means \mathbf{y} , i.e.

$$\mathbf{Y} \sim \text{Poisson}(\mathbf{y}). \quad (1.21)$$

Following a similar derivation given in Harmany *et al.* (2012), we can now write:

$$p(\mathbf{Y}|\mathbf{y}) = \prod_{i=1}^{N_{fp}} \frac{(e_i^T \mathbf{y})^{Y_i}}{Y_i!} \exp(-e_i^T \mathbf{y}) \quad (1.22)$$

where e_i is the unit vector of the i -th canonical basis. Substituting for \mathbf{y} , we obtain the negative log-likelihood as:

$$F(\mathbf{y}) = \mathbf{1}_{N_{fp}}^T \mathbf{y} - \sum_{i=1}^{N_{fp}} Y_i \log(e_i^T \mathbf{y}) + C \quad (1.23)$$

where $\mathbf{1}$ is an all-ones vector of size N_{fp} , and C is a constant independent of $\hat{\mathbf{v}}$. Let $\tilde{\mathbf{A}} := \frac{2n}{N_{fp}^2} \mathbf{A}$. Then the gradient with respect to $\hat{\mathbf{v}}$ is given by:

$$\nabla_{\hat{\mathbf{v}}} F(\mathbf{y}) = \tilde{\mathbf{A}}^T \mathbf{1} - \sum_{i=1}^{N_{fp}} \frac{Y_i}{e_i^T \mathbf{y}} \tilde{\mathbf{A}}^T e_i \quad (1.24)$$

Therefore the Hessian is:

$$\nabla_{\hat{\mathbf{v}}}^2 F(\mathbf{y}) = \tilde{\mathbf{A}}^T \sum_{i=1}^{N_{fp}} \frac{Y_i}{(e_i^T \mathbf{y})^2} e_i e_i^T \tilde{\mathbf{A}} \quad (1.25)$$

To obtain the Fisher information matrix $\mathbf{I}(\hat{\mathbf{v}})$, we take the negative-expectation of this quantity:

$$\mathbf{I}(\hat{\mathbf{v}}) = -E[\nabla_{\hat{\mathbf{v}}}^2 F(\mathbf{y})] = -\tilde{\mathbf{A}}^T \sum_{i=1}^{N_{fp}} \frac{E[Y_i]}{(e_i^T \mathbf{y})^2} e_i e_i^T \tilde{\mathbf{A}} \quad (1.26)$$

But since $E[Y_i] = e_i^T \mathbf{y}$, this simplifies to:

$$\mathbf{I}(\hat{\mathbf{v}}) = \tilde{\mathbf{A}}^T \mathbf{D} \tilde{\mathbf{A}} \quad (1.27)$$

where \mathbf{D} is a diagonal matrix with $\mathbf{D}_{ii} = \frac{1}{(e_i^T \mathbf{y})}$.

Therefore the Cramer Rao Lower Bound (CRLB) for the variance of each estimated Fourier coefficient is given by:

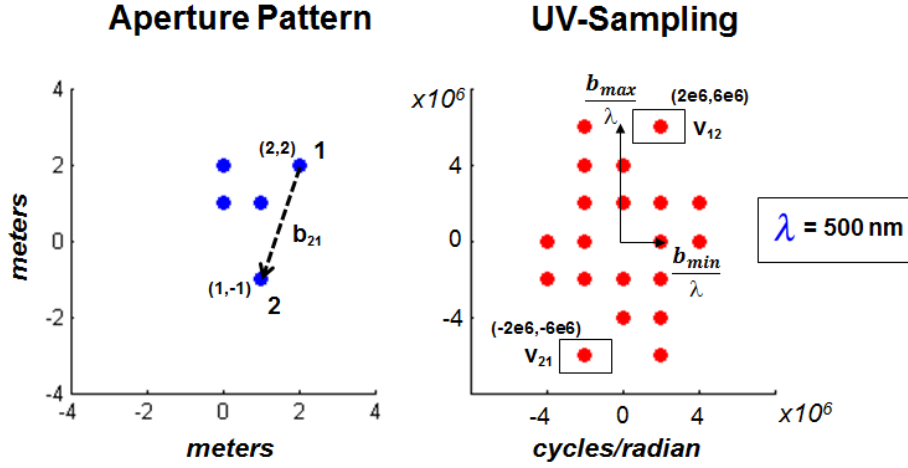


Figure 1.4: The relationship between aperture patterns and Fourier sampling

$$\text{var}(\hat{v}_i) \geq [\mathbf{I}(\hat{\mathbf{v}})^{-1}]_{ii} \quad (1.28)$$

This expression matches the result given by Zmuidzinas (2003).

1.5 Key Parameters of an Interferometer: Resolution, Field-of-View, and Undersampling Ratio

In the Section 1.3 we established that the interference fringe generated by a pair of apertures separated by a vector \mathbf{b} encodes the amplitude and phase of the Fourier Transform of the scene at spatial frequency $\frac{\mathbf{b}}{\lambda}$. For real images each baseline actually contributes two spatial frequency samples, as the Fourier Transform of real images is conjugate symmetric. As an example, consider the aperture pattern in the left panel of Figure 1.4. Suppose we are interfering light at a wavelength of 500 nm. The baseline \mathbf{b}_{12} formed by apertures 1 and 2 generates a fringe encoding the complex visibility \mathbf{v}_{12} at a spatial frequency of $\frac{\mathbf{b}_{12}}{\lambda} = (2e6, 6e6)$ cycles/radian, and its conjugate \mathbf{v}_{12}^* at a spatial frequency $-(2e6, 6e6)$ cycles/radian.

Consider the Cartesian frequency sampling grid with gridpoint spacing Δ_{min} as shown

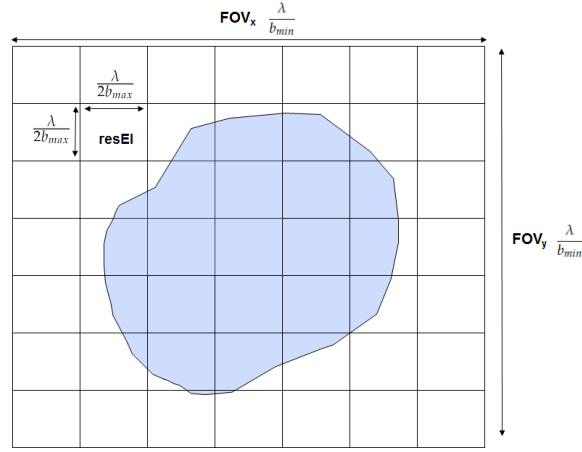


Figure 1.5: Representation of a scene as a Field-of-View comprised of resolution elements

in Figure. Suppose we sample a 2D-signal in the Fourier domain with sampling interval $\Delta_{s,freq} = \frac{1}{2R}$ along each dimension. By the well-known Nyquist Sampling Theorem, we know that only images spatially-limited to $\pm R$ along each dimension can be uniquely represented with these samples; images beyond this extent will suffer from aliasing. Therefore an interferometer whose sampled spatial frequencies lie on a grid with spacing $\Delta_{min} = \frac{b_{min}}{\lambda}$ has an unambiguous Field-of-View area of $\frac{1}{\Delta_{min}^2} = \frac{\lambda^2}{b_{min}^2}$. Conversely, by the space-frequency duality of the Nyquist Theorem, we know that we can uniquely represent an image bandlimited to frequency extent $\pm L$ by sampling in space at an interval of $\Delta_{s,space} = \frac{1}{2L}$. Given the maximum spatial frequency observed by our interferometer is $L = \Delta_{max} = \frac{b_{max}}{\lambda}$, the bandlimited approximation of the image is uniquely defined by samples spaced $\frac{\lambda}{2b_{max}}$ apart along each dimension. Therefore the size of the smallest resolvable element in the image, which we call the *resolution element*, is given by $\frac{\lambda^2}{4b_{max}^2}$. Figure 1.5 depicts the notions of Field-of-View (FOV) and resolution element (resEl) visually.

With the areas of the Field-of-View and resolution element in hand, we can now compute the number of resolution elements in the bandlimited approximation of the image as the ratio of these two quantities, or:

$$N_{resEls} = 4r^2 \quad (1.29)$$

where $r = \frac{b_{max}}{b_{min}}$.

We now define our undersampling ρ as the ratio of the number of distinct Fourier samples to the number of resolution elements N_{resEls} in the image. Each Fourier sample we measure is actually a pair of samples since for real images, the value of the Fourier Transform at a given spatial frequency is the conjugate of that at the corresponding negative spatial frequency. Hence we have:

$$\rho = \frac{2 \binom{N_{ap}}{2}}{N_{resEls}} = \frac{1}{2r^2} \binom{N_{ap}}{2} \quad (1.30)$$

1.6 Atmospheric and Instrumental Phase Noise

Our analysis thus far has assumed that the effective path between the beam-combiner and the target depends merely on the source-instrument geometry depicted in Figure 1.2. In practice there are several deterministic and random sources of path variation across the aperture array. Common deterministic sources of path variation include that in the optical paths traversed between the instrument's front-end and beam-combiner. The predominant random sources of delay are atmospheric turbulence and instrumental vibration. Atmospheric turbulence, which is typically the more stressing of these two sources, alters the optical path traversed by the wave arriving at each aperture in a nonuniform and time-varying manner.

Let us consider the total path alteration due to all of the sources mentioned. Let us then define the corresponding phase shift resulting from this alteration at aperture j as $e^{i\phi_{j,total}}$. If these phase shifts are carried through the analysis in Section 1.3, we obtain:

$$\langle \mathbf{y}(\boldsymbol{\rho}) \rangle = \frac{nN_{ap}}{N_{fp}^2} + \frac{2n}{N_{fp}^2} \sum_{(j,k)}^{\binom{N_{ap}}{2}} \gamma_{jk} \cos(\boldsymbol{\rho} \cdot \Delta \mathbf{r}_{jk} + \theta_{jk} + \phi_{jk}) \quad (1.31)$$

where $\phi_{jk} := \phi_{j,total} - \phi_{k,total}$ is the difference between the total phase shifts at the two apertures in the baseline associated with aperture pair (i, j) .

Let us now examine a traditional method for eliminating the effect of this phase noise which was first suggested by Jennison (1958) in the context of inteferometry at radio wave-

lengths. Suppose the atmosphere adds phases ϕ_j and ϕ_k to apertures k and j , respectively. The result is that the observed complex visibility is given by:

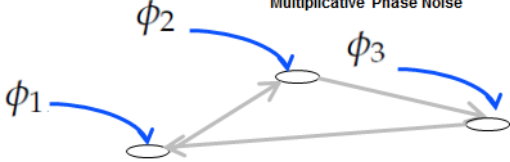
$$\tilde{\mathbf{v}}_{jk} = \mathbf{v}_{jk} e^{i(\phi_j - \phi_k)} \quad (1.32)$$

We can eliminate such nuisance factors by forming a triple product \mathbf{g} of the Fourier phasors along a triangle of baselines (e.g. $\mathbf{g}_{123} := \tilde{\mathbf{v}}_{12}\tilde{\mathbf{v}}_{23}\tilde{\mathbf{v}}_{31}$). As shown in Figure 1.6, this special triple product (known as the *bispectrum*) cancels the atmospheric phase terms, leaving only the desired Fourier information (i.e. the $\{\theta\}$). The phase of the bispectrum $\angle \mathbf{g}$, which is known as the *closure phase*, can hence be used to recover estimates for the Fourier phases. Consider an interferometer that measures all baselines among N_{ap} apertures. Out of the $\binom{N_{ap}}{3}$ possible triangles, only $\binom{N_{ap}-1}{2}$ of the associated closure phase relations are linearly-independent (Readhead *et al.*, 1988). If we combine these closure phases with the Fourier magnitude estimates for each of the $\binom{N_{ap}}{2}$ baselines, we now have a set of atmospheric-invariant observables with which we can attempt image reconstruction.

The Atmosphere Problem

$$\tilde{\mathbf{v}}_{jk} = \mathbf{v}_{jk} e^{i(\phi_j - \phi_k)}$$

Multiplicative Phase Noise



Fourier triple products (*bispectra*) cancel atmosphere

$$\mathbf{g}_{123} = \tilde{\mathbf{v}}_{12} \tilde{\mathbf{v}}_{23} \tilde{\mathbf{v}}_{31}$$

$$\angle \tilde{\mathbf{v}}_{12} = \theta_{12} + \cancel{\phi_1} - \cancel{\phi_2}$$

$$\angle \tilde{\mathbf{v}}_{23} = \theta_{23} + \cancel{\phi_2} - \cancel{\phi_3}$$

$$+ \angle \tilde{\mathbf{v}}_{31} = \theta_{31} + \cancel{\phi_3} - \cancel{\phi_1}$$

$$\angle \mathbf{g}_{123} = \theta_{12} + \theta_{23} + \theta_{31}$$

Figure 1.6: *Eliminating the effect of atmospheric turbulence with phase closure*

Chapter 2

Leveraging compressed sensing techniques in optical interferometry

2.1 Citation to Previously-Published Work

This chapter contains text and figures published previously in the following paper:

Kurien, B. G., Rachlin, Y., Shah, V. N., Ashcom, J. B. and Tarokh, V. (2014). Compressed sensing techniques for image reconstruction in optical interferometry. In Imaging and Applied Optics 2014, Optical Society of America, p. SM2F.3.

2.2 Chapter Overview

Optical interferometers image a scene by sampling the spatial frequencies that comprise its 2D Fourier Transform. The sample set is typically much smaller than the number of resolvable elements in the instrument's field-of-view (FOV). In the related field of radio interferometry, sparse recovery (SR) techniques based on recent seminal work (see, e.g., Wiaux *et al.* (2009)) have proven successful in regularizing the ill-posed problem arising from this undersampling. In contrast, atmospheric turbulence precludes the availability of direct Fourier phase information in optical interferometry, and hence CS techniques do not

directly apply. We have developed and validated a robust algorithmic interface between the Fourier magnitude and bispectrum observables available in the optical interferometry and SR regularizations known as Total Variation Minimization for which fast solvers exist.

2.3 Problem Statement and Approach

There are two principal challenges that one encounters when trying to reconstruct images from optical interferometric measurements. Each aperture pair (or baseline) in an interferometer samples the Fourier transform of the image in its field of view at a spatial frequency $\frac{\mathbf{b}}{\lambda}$, where \mathbf{b} is the length of the baseline, and λ is the wavelength of light collected. Interferometers rarely measure enough baselines to fully sample the Fourier transform. Hence reconstruction of the image from such undersampled measurement set is an ill-posed (or underdetermined) problem; we must recover intensities for N_{resEl} resolution elements in the scene under observation from $m \ll N_{resEl}$ spatial frequency measurements corresponding to the available baseline pairs in the telescope array. To make the problem well-posed, additional constraints must be enforced. Traditionally regularization approaches based on the so-called *CLEAN* algorithm due to Högbom (1974), which implicitly fits the measurements to an image model consisting of a sparse collection of point-like sources, have been used to solve this problem.

It is well-known in imaging that successful regularization schemes apply constraints which accurately capture the properties of the image we wish to reconstruct. Likewise the success of the CLEAN algorithm has been largely limited to astronomical scenes which closely match its point-source-collection assumption. The astronomical community has hence been led to consider more generally-applicable prior models. A property shared by an overwhelmingly-large fraction of natural images is *compressibility*. Compressible images are those which can be well-approximated by a sparse representation in some domain. The smallness of an image's L1-norm¹, as applied to either directly to the image pixels, their

¹The L1-norm of a vectorized image \mathbf{x} , denoted $\|\mathbf{x}\|_1$, is the sum of the absolute values of the entries in \mathbf{x}

gradient, or their wavelet coefficients, has proven to be a remarkably effective proxy for the image's compressibility. This observation and techniques that exploit it are at the core of the field of *sparse recovery* (SR) as well as the tightly-linked field of *compressed sensing* (CS) (Candès *et al.*, 2006) (Donoho, 2006). Though a wide variety of sparse-recovery techniques have arisen in the past few decades, the vast majority of them regularize problems with a linear observation model:

$$\mathbf{y} = \mathbf{F}\mathbf{x} + \mathbf{n} \quad (2.1)$$

where \mathbf{y} is the measurement vector, \mathbf{x} is a vector representing the unknown signal or image, \mathbf{F} is the measurement matrix, and \mathbf{n} represents additive measurement noise.

While sparse-recovery techniques have been applied successfully to reconstruction in radio interferometry (Wiaux *et al.*, 2009), optical interferometry, on the other hand, poses additional challenges beyond the Fourier-undersampling problem discussed above. Namely, atmospheric turbulence necessitates a non-linear formulation of the reconstruction problem. Recall from Chapter 1 that there are $\binom{N_{ap}}{2}$ unknown Fourier phases as well as $N_{ap} - 1$ unknown piston differences, and hence inference of the Fourier phases from the $\binom{N_{ap}}{2}$ phase measurements available in a non-redundant array is inherently ill-posed. The traditional means of mitigating this issue has been to again impose prior constraints on the reconstructed image (e.g. constraints on the scene's compactness, sparsity, or smoothness). In particular, a myriad of so-called *self-calibration* algorithms have been developed (see, e.g., Pearson and Readhead (1984)) which alternate between estimation of the Fourier phases and estimation of the image coefficients subject to the imposed constraints. Such algorithms have been shown to be successful in high-flux scenarios in which the interferometric measurements have a high signal-to-noise ratio (SNR). In low-flux scenarios on the other hand, one must integrate observables over a long time period in order to build sufficient SNR for reliable image reconstruction. Since integration beyond the coherence time of the atmosphere would result in fringe blurring, one is then led to the formation and subsequent integration of atmosphere-invariant observables from each atmosphere-coherence time (or

frame).

In Section 1.6 we introduced the bispectrum observable as the classical basis for atmosphere-invariant inference in optical interferometry. Formed as the product of three fringe phasors associated with the sides of a baseline triangle (e.g. \mathbf{b}_{12} , \mathbf{b}_{23} , and \mathbf{b}_{31}), the bispectrum is a non-linear function of the complex visibilities of the image. Recall that the atmosphere-induced terms cancel in these products and hence, like the Fourier magnitudes, these so-called *bispectra* are atmosphere-invariant observables. However, for a non-redundant array with $\binom{N}{2}$ distinct baselines, recovery of the Fourier phases from the bispectra phases (i.e. the *closure phases*) remains ill-posed since there are only $\binom{N-1}{2}$ independent closure phases (Readhead *et al.*, 1988). Successful bispectra-based image reconstruction remains feasible in spite of this ill-posedness (see e.g. Thiébaud (2013), Besnerais *et al.* (2008)), but again prior constraints (e.g. on the image support) must be enforced to regularize the reconstruction.

An important property of the bispectrum which we will exploit in our reconstruction method is its invariance to scene translation. To illustrate this property, let us denote the spatial frequency vectors of two sides of a bispectrum triangle as \mathbf{u} and \mathbf{v} , respectively. Then the spatial frequency of the remaining side will be $\mathbf{w} := -(\mathbf{u} + \mathbf{v})$. If $\hat{g}(\mathbf{p})$ is the corresponding (normalized) bispectrum of a scene at its original position \mathbf{p} , then by the shift property of the Fourier transform, the bispectrum of the translated image becomes:

$$\hat{g}(\mathbf{p} + \delta\mathbf{p}) = e^{j(\theta_{\mathbf{u}} + \mathbf{u} \cdot \delta\mathbf{p})} e^{j(\theta_{\mathbf{v}} + \mathbf{v} \cdot \delta\mathbf{p})} e^{j(\theta_{-(\mathbf{u}+\mathbf{v})} - (\mathbf{u}+\mathbf{v}) \cdot \delta\mathbf{p})} \quad (2.2)$$

$$= e^{j(\theta_{\mathbf{u}} + \theta_{\mathbf{v}} + \theta_{-(\mathbf{u}+\mathbf{v})})} = \hat{g}(\mathbf{p}) \quad (2.3)$$

where we have dropped the aperture subscripts for purposes of generality.

Since the bispectra are non-linear functions of their underlying image, reconstruction from these observables does not map directly onto the linear SR framework in Equation (2.1). To address this challenge, we developed an interface which first recovers Fourier component estimates from bispectra and Fourier magnitudes, using a nonlinear least squares solver.

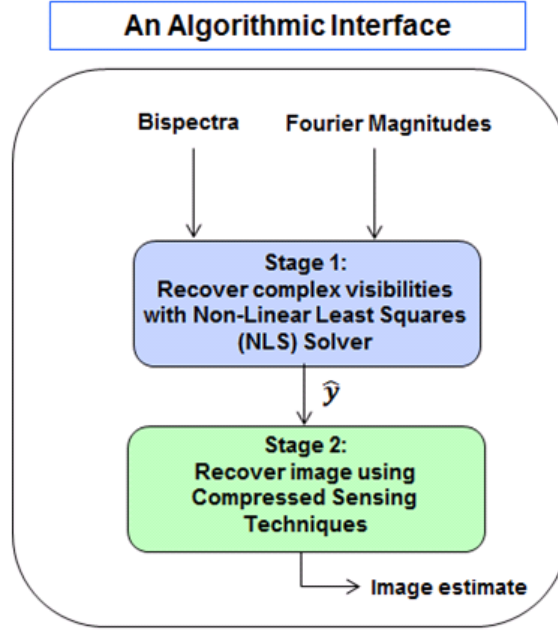


Figure 2.1: *Overview of Proposed Two-Stage Approach*

With the Fourier component estimates at hand, the image estimate can then be recovered using standard SR techniques. The approach is diagrammed in Figure 2.1 below.

2.4 Description of Stage 1 of Algorithm

The goal of Stage 1 of our algorithm is to use the measured bispectra and Fourier magnitudes to recover estimates of the true Fourier components of the image. Recall from Section 1.3 that the fringes provide measurements of the Fourier magnitudes directly. Hence it remains to recover the Fourier phase information. Since there are $\binom{N_{ap}}{2}$ unknown Fourier phases in a non-redundant array but only $\binom{N_{ap}-1}{2}$ independent closure phases, this sub-problem is itself ill-posed. Namely, we cannot hope to find a unique solution for the Fourier phases simply by minimizing the residuals of the bispectra data. Instead our algorithm employs a gradient-based search to find the phase set that minimizes a regularized objective consisting of a data term and a prior term. The data term is the sum of the squared bispectrum residuals. The prior term is a metric that favors those Fourier phase vectors corresponding

to a spatially-compact intensity profile in the image domain. This kind of joint metric has been suggested for optical interferometric applications before in the work of Thiébaud (2013). However, whereas the method in Thiébaud (2013) searches for a metric-minimizing image, our method first searches for the metric-minimizing baseline phase set, which is inherently a smaller set given the undersampled nature of the problem.

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \mathbf{f}_{data}(\boldsymbol{\theta}, \bar{\mathbf{g}}, \mathbf{a}) + \mu \mathbf{f}_{prior}(\boldsymbol{\theta}, \mathbf{a}) \quad (2.4)$$

The data (or *residual*) and prior terms are defined as follows.

$$\mathbf{f}_{data}(\boldsymbol{\theta}, \bar{\mathbf{g}}, \mathbf{a}) = \sum_{k=1}^{N_b} \left\| \frac{1}{w_k} (\bar{g}_k - a_{k1}a_{k2}a_{k3} \exp(j(\theta_{k1} + \theta_{k2} + \theta_{k3}))) \right\|^2 \quad (2.5)$$

where \bar{g}_k is the k -th bispectrum measurement, $\{\theta_{ki}\}$ are the unknown Fourier phases in the k -th bispectrum, the $\{a_{ki}\}$ are estimates for the magnitudes of the Fourier coefficients in the k -th bispectrum, and w_k is a weighting factor proportional to the estimated variance of the k -th bispectrum.

$$\mathbf{f}_{prior}(\boldsymbol{\theta}, \mathbf{a}) = \left\| \text{Re}[\mathbf{F}^* \mathbf{y}_{\mathbf{a}, \boldsymbol{\theta}}] \diamond (\mathbf{1} - \mathbf{h}_{gaussian}) \right\|^2 \quad (2.6)$$

where \mathbf{F} is the partial DFT matrix whose rows are 2D-sinusoidal basis functions sampled by the array, $\mathbf{y}_{\mathbf{a}, \boldsymbol{\theta}}$ is the vector of estimated Fourier coefficients with magnitudes given by \mathbf{a} and phases given by $\boldsymbol{\theta}$, $\mathbf{h}_{gaussian}$ denotes a centered Gaussian window with a peak value of unity, $\mathbf{1}$ denotes the vector whose entries are all 1, and the operation \diamond denotes point-wise multiplication.

An intuitive interpretation of the regularization above can be obtained by decomposing the computation into steps. We seek to fit the bispectra with a set of Fourier phases $\hat{\boldsymbol{\theta}}$ which together with the estimated Fourier magnitudes produce an image with compact energy. The data term in the metric \mathbf{f}_{data} is a sum of the squared residuals of measured bispectra with respect to the bispectra associated with the phase vector $\hat{\boldsymbol{\theta}}$. The steps for computation of the prior term are as follows. First we associate the phases in $\boldsymbol{\theta}$ with their corresponding

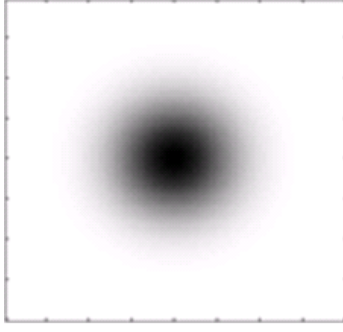


Figure 2.2: *Inverted Gaussian penalty function ($\mathbb{1} - \mathbf{h}_{\text{gaussian}}$); the dark region represents an area of near-zero penalty*

magnitude estimates in \mathbf{a} to form estimates for the complex Fourier components $\mathbf{y}_{\mathbf{a},\theta}$. We then apply the pseudo-inverse of the partial DFT matrix \mathbf{F} to the vector $\mathbf{y}_{\mathbf{a},\theta}$ and retain the real part ², thereby obtaining an image-space representation of $\mathbf{y}_{\mathbf{a},\theta}$. Finally, we perform point-wise multiplication of this image with a function of the form shown in Figure and compute the norm of the result. This operation computes a compactness metric which penalizes $\mathbf{y}_{\mathbf{a},\theta}$ in proportion to its spatial energy spread. Note that the actual position of the Gaussian taper is immaterial; as established above, the bispectrum is a translation-invariant quantity and hence the joint metric will be invariant to shifts of the same object.

The joint data-prior metric in Equation (2.4) can be minimized using a suitable non-linear least-squares solver. For our processing, we have elected to use the MATLAB[®] implementation (MATLAB, 2012) of the trust-region reflective non-linear least-squares solver.

2.5 Description of Stage 2 of Algorithm

Stage 2 takes the Fourier estimates obtained in Stage 1 and attempts to recover an image using fast sparse-recovery (SR) techniques. This amounts to solving the linear, under-determined inference problem:

²Since the rows of \mathbf{F} are orthogonal, the pseudo-inverse is obtained by a Hermitian transpose operation \mathbf{F}^*

$$\hat{\mathbf{y}} = \tilde{\mathbf{F}}\mathbf{x} + \mathbf{n} \quad (2.7)$$

where $\hat{\mathbf{y}}$ is the vector of Fourier estimates from Stage 1, $\tilde{\mathbf{F}}$ is the partial-DFT matrix whose rows are the complex-sinusoidal basis functions corresponding to the spatial frequencies sampled by the array³, and \mathbf{n} is the (unknown) residual error in these estimates. Since in practice there are fewer available Fourier measurements than the desired number of resolution elements in the image, this inference problem is ill-posed from a Nyquist-sampling perspective. Hence once again we must regularize the problem by imposing constraints on the image. We selected two different SR regularization strategies due to their widespread success in solving similar inverse problems involving Fourier undersampling. The first of these approaches, Basis Pursuit Denoising (BPDN), seeks to find the image of smallest L1-norm which also agrees with the estimated Fourier component estimates to within a certain tolerance ϵ . We performed this L1-minimization in the wavelet domain as opposed to the pixel domain, as it is well-known that natural images tend to be more compressible in the former domain than in the latter. The other regularization we tested, Total Variation Minimization (TV-Min), is very similar, but seeks instead to minimize another quantity known to be compressible in natural images: the L1-norm of the gradient of the image. The two regularization techniques are specified mathematically below:

$$\text{(Basis Pursuit Denoising): } \hat{\mathbf{x}} = \arg \min_{\alpha} \|\alpha\|_1 \text{ subject to: } \|\mathbf{F}\Psi^{-1}\alpha - \hat{\mathbf{y}}\|_2 \leq \epsilon \quad (2.8)$$

$$\text{(Total Variation Minimization): } \hat{\mathbf{x}} = \arg \min_{\alpha} \|\alpha\|_{TV} \text{ subject to: } \|\mathbf{F}\alpha - \hat{\mathbf{y}}\|_2 \leq \epsilon \quad (2.9)$$

where $\|\alpha\|_{TV}$ is the L1-norm of the 2D pixel gradient of the image, i.e. $\|\alpha\|_{TV} := \sum_{i,j} \|\nabla \alpha[i,j]\|_1$.

³Note that $\tilde{\mathbf{F}}$ need not be the same as \mathbf{F} , as we are free to assume an independent pixel resolution in each case. This resolution then specifies the granularity of the sampling of the sinusoids forming the rows of the matrix.

The software package NESTA (Becker *et al.*, 2011) was used to perform the Stage 2 optimizations above.

2.6 Algorithm Performance

2.6.1 Laboratory Validation

To test the effectiveness of our algorithm when applied to real interferometric data, we collected fringe data with MIT Lincoln Laboratory’s (MIT/LL) Fizeau Interferometer. The interferometer is a reconfigurable fiber-coupled system that allows all $\binom{20}{2}$ baseline conjugate pairs to be measured simultaneously by projecting light from all apertures onto a Fizeau beam-combiner. For our experiments, we used a 20-aperture compact, non-redundant pattern of the Golay type (Golay, 1970). The pattern is shown in Figure 2.3. The ratio r of the maximum baseline to the minimum baseline was 17 in both x and y spatial coordinates. Hence the number of resolution elements required to uniquely specify our diffraction-limited scene was $(2r)^2$, yielding an undersampling ratio of $\rho = \frac{2 \cdot 190}{(2 \cdot 17)^2} \approx 0.33$. The interferometer was illuminated with the far-field projection of a range of targets: chrome-on-glass transparency masks were photolithographically-prepared and placed at the focus of an $\frac{f}{12}$ three-mirror off-axis telescope. When the target is illuminated with the white light, the far-field projection of the target is produced at the telescope aperture, where it can be sampled. As a reference for reconstruction, we consider a reduction of the actual target to the fundamental diffraction-limited resolution of the instrument as shown in Figure 2.4.

5000-frame reconstruction results are shown in Figure 2.5 for the following photoelectron (pe) levels: (from left to right) $320e3$ pe/aperture/frame, $54e3$ pe/aperture/frame, and $15e3$ pe/aperture/frame.

2.6.2 Simulation

In addition to the Laboratory validation, we also conducted 5000-frame simulations with the same aperture pattern to compare algorithm performance in the presence of unknown

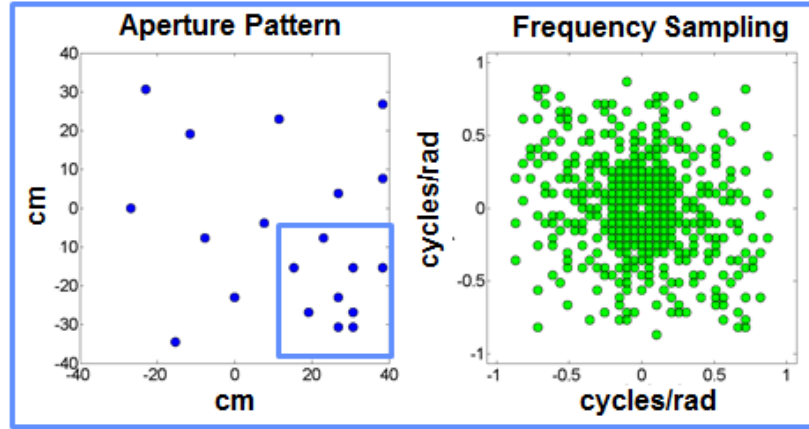


Figure 2.3: Non-redundant Golay 20-aperture pattern (left) and corresponding UV-sampling (right)

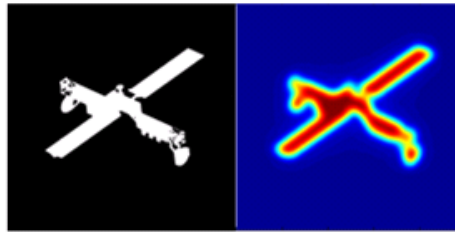


Figure 2.4: Target chrome mask (left) and corresponding truth image at diffraction-limited resolution (right)

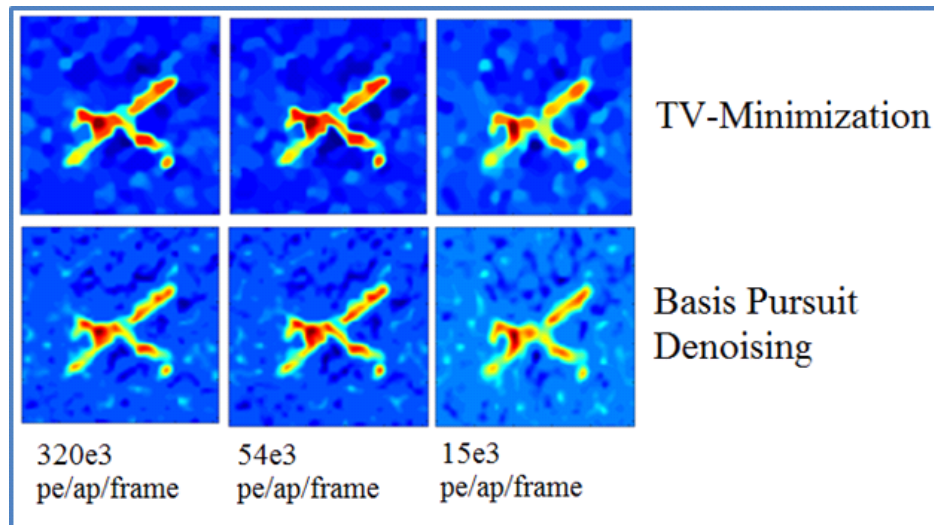


Figure 2.5: Image Reconstruction Results from Laboratory Validation

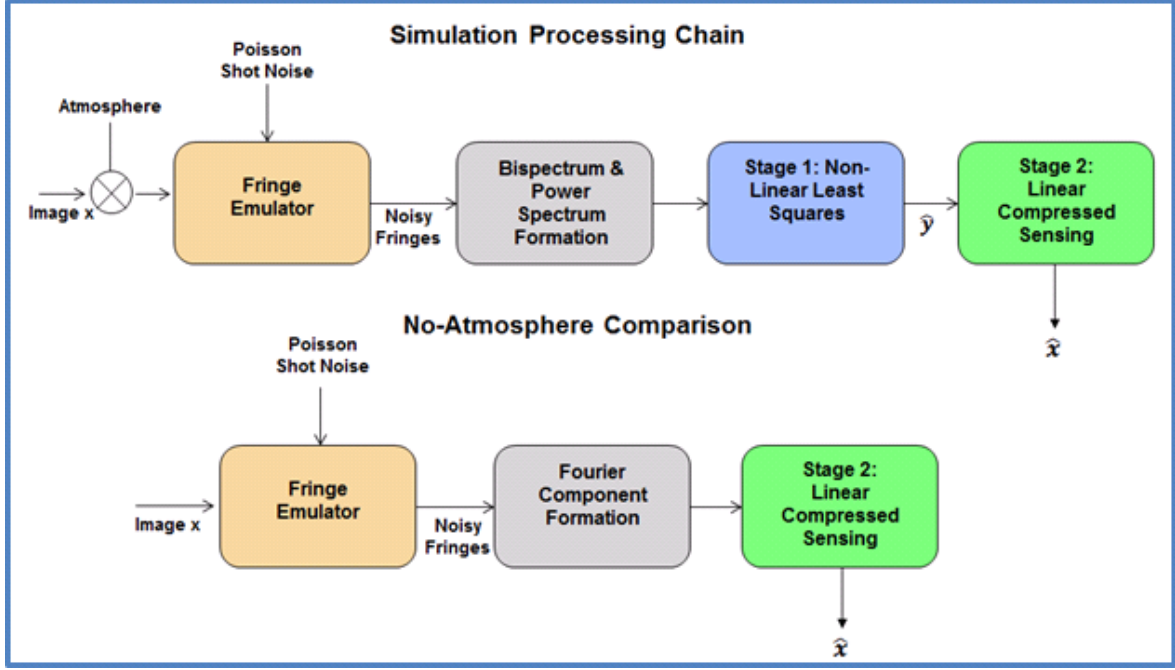


Figure 2.6: Image Reconstruction Results from Simulations

atmospheric distortion against that if this distortion were known. Note that in the latter *atmosphere-oracle* case, bispectrum formation is not required and direct SR reconstruction from complex amplitudes is feasible. Hence this comparison allows a quantitative measure of the reconstruction penalty suffered from bispectrum formation. Our comparative simulation scheme is depicted in the block diagram in Figure 2.6.

We used two Image Quality Metrics to assess performance: the standard metric Normalized Mean-Squared Error (NMSE) and the Structural Similarity (SSIM) metric (Wang *et al.*, 2004). The NMSE is defined as:

$$NMSE = 20 \log \frac{\|\hat{\mathbf{x}} - \mathbf{x}_0\|}{\mathbf{x}_0} \quad (2.10)$$

Neither metric captured the visual similarity of the raw reconstructions of our two-stage algorithm relative to ground truth. To address this, we uniformly hard-thresholded each of the raw reconstruction pixels before evaluating the metrics. Figure 2.7 shows our two-stage reconstructions before thresholding (left), the same reconstructions after thresholding

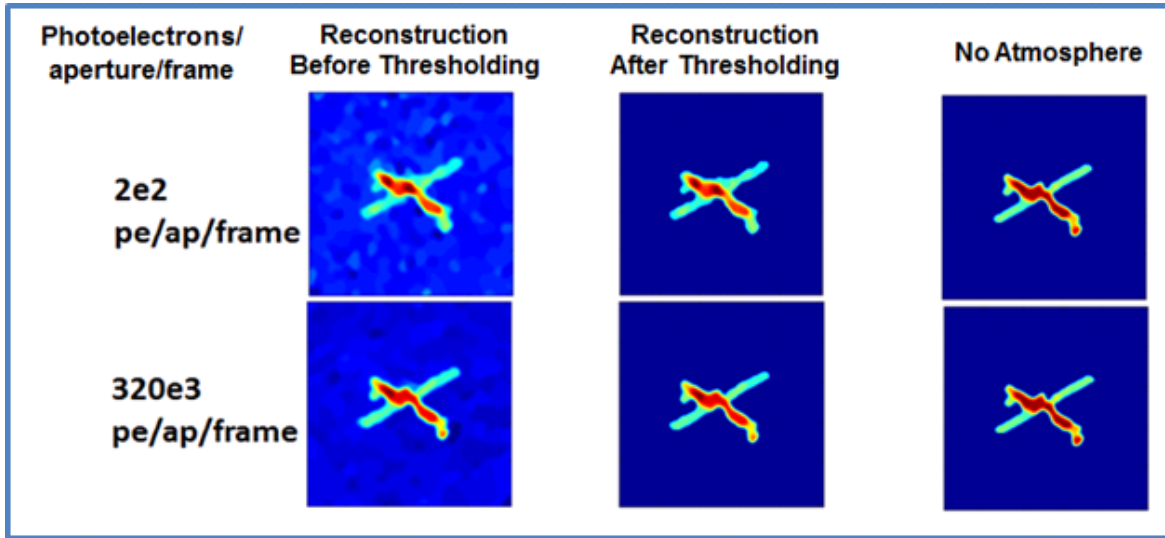


Figure 2.7: Image Reconstruction Results from Simulations

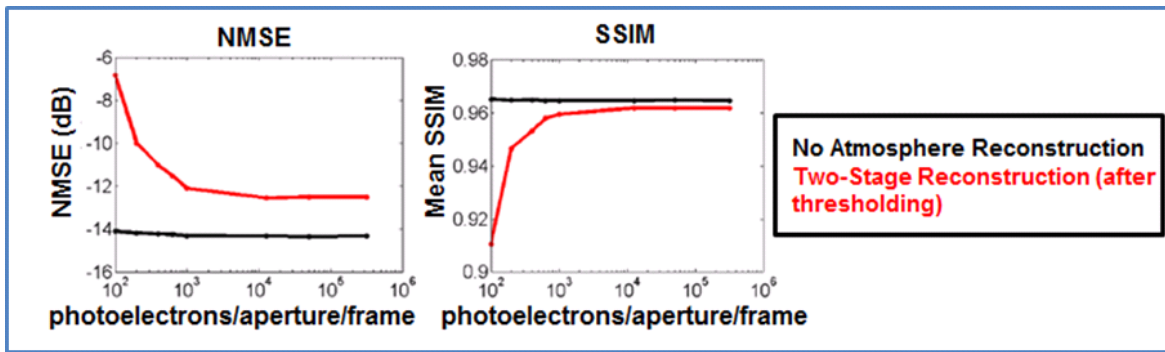


Figure 2.8: Algorithm Performance in Simulation

(middle), and the atmosphere-oracle reconstructions (right). Results from both a low-light-level scenario (top row) and a high-light-level scenario (bottom row) are shown. Figure 2.8 uses NMSE and SSIM image quality metrics to quantify reconstruction performance for our two-stage-algorithm reconstructions after thresholding (red), and the atmosphere-oracle reconstruction (black). We see that the algorithm achieves performance close to that of the atmosphere oracle for a wide range of light levels.

Chapter 3

Pattern design criteria for uniqueness in phase recovery

3.1 Citation to Work under Review

This chapter contains text and figures submitted for publication to the Monthly Notices of the Royal Astronomical Society (MNRAS).

Kurien, B., Tarokh, V., Ashcom, J., Rachlin, Y. and Shah, V. (2016). Resolving phase ambiguities in the calibration of redundant interferometric arrays: implications for array design. Monthly Notices of the Royal Astronomical Society, (submitted, March 4, 2016)

3.2 Chapter Overview

We provide new results enabling robust interferometric image reconstruction in the presence of unknown aperture piston variation via the technique of Redundant Spacing Calibration (RSC). The RSC technique uses redundant measurements of the same interferometric baseline with different pairs of apertures to reveal the piston variation among these pairs. In both optical and radio interferometry, the presence of phase-wrapping ambiguities in the measurements is a fundamental issue that needs to be addressed for reliable image

reconstruction. In this paper, we show that these ambiguities affect recently-developed RSC phasor-based reconstruction approaches operating on the complex visibilities, as well as traditional phase-based approaches operating on their logarithm. We also derive new sufficient conditions for an interferometric array to be immune to these ambiguities in the sense that their effect can be rendered benign in image reconstruction. This property, which we call *wrap-invariance*, has implications for the reliability of imaging via classical three-baseline phase closures as well as generalized closures. We show that wrap-invariance is conferred upon arrays whose interferometric graph satisfies a certain cycle-free condition. For cases in which this condition is not satisfied, a simple algorithm is provided for identifying those graph cycles which prevent its satisfaction. We apply this algorithm to diagnose and correct a member of a pattern family popular in the literature.

Before we discuss phase-wrapping ambiguities, we first review a few key preliminary mathematical notions from lattice theory in the next Section.

3.3 Preliminaries

3.3.1 Lattices

A *lattice* is a mathematical object describing a repeating pattern of discrete points in space. It arises in a variety of scientific and engineering contexts including crystallography and communication theory. We will define a lattice Λ as a linear combination of vectors in which the coefficients are integers, i.e.

$$\Lambda = \left\{ \sum_{i=1}^m a_i \mathbf{v}_i \mid \forall a_i \in \mathbb{Z} \right\} \quad (3.1)$$

If the generating vectors \mathbf{v}_i are not linearly-independent, one might ask for another, more compact description of Λ . Consider one such candidate set of linearly-independent vectors \mathbf{b}_j . If all integer linear combinations of \mathbf{b}_j lie in Λ , and every lattice point in Λ can be represented as an integer linear combination of the \mathbf{b}_j , then we say that the vectors \mathbf{v}_j form a *basis* for the lattice Λ .

Many lattice algorithms, including algorithms for the Closest Vector Problem introduced in the next section, require a *reduced* lattice basis. A *reduced* basis is one consisting of vectors which are short and nearly-orthogonal. A well-known procedure to create a basis for a given lattice with short, nearly-orthogonal vectors is known as the *LLL* algorithm due to Lenstra *et al.* (1982).

An LLL-based routine for forming a reduced lattice basis from a set of generating vectors which are not necessarily linearly-independent is given in Cohen (1993).

3.3.2 The Closest-Vector-Problem

A well-known problem in lattice theory is the *Closest-Vector-Problem*, which can be described as follows:

Problem 3.3.1. *Given a basis $\{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n\}$ for a lattice in \mathbb{R}^m and a vector $\mathbf{w} \in \mathbb{R}^m$, find the point in the lattice closest in Euclidean distance from \mathbf{w} .*

Several algorithms exist for solving this problem. A popular class of algorithms, known as the Sphere-Decoding algorithms, are efficient searches for the closest lattice point within a hypersphere of a certain radius centered on the input vector (see e.g. Agrell *et al.* (2002)). For the simulations in this chapter, we instead use the lower-complexity Babai Nearest Plane (Babai-NP) algorithm (Babai, 1986). For lattice bases which are nearly orthogonal, this algorithm offers reliable, albeit not guaranteed, performance in practice. Pseudo-code for one implementation of this algorithm due to Galbraith (2012) is given in listing Algorithm 1.

Algorithm 1 Babai Nearest-Plane Algorithm

Input: Basis for lattice Λ (i.e. $\{\mathbf{b}_1, \dots, \mathbf{b}_n\}$), and $\mathbf{w} \in \mathbb{R}^m$

Output: Element \mathbf{v}^* nearest to \mathbf{w} in Λ

 Compute orthogonal basis $\{\mathbf{b}_1^*, \dots, \mathbf{b}_n^*\}$ using Gram-Schmidt procedure

for $i = n$ **downto** 1 **do**

 Set $l_i = \frac{\langle \mathbf{w}_i, \mathbf{b}_i^* \rangle}{\langle \mathbf{b}_i^*, \mathbf{b}_i^* \rangle}$

 Set $\mathbf{y}_i = \lfloor l_i \rfloor \mathbf{b}_i$

 Set $\mathbf{w}_{i-1} = \mathbf{w}_i - (l_i - \lfloor l_i \rfloor) \mathbf{b}_i^* - \lfloor l_i \rfloor \mathbf{b}_i$

end for

return $\mathbf{v} = \mathbf{y}_1 + \dots + \mathbf{y}_n$

A conceptual interpretation of this recursive algorithm is as follows. At the root level, the algorithm finds the vector in Λ closest to \mathbf{w} ; in other words it solves the CVP problem (Λ_n, \mathbf{w}) . First consider the subspace U spanned by the first $n - 1$ basis vectors, i.e. $U = \text{span}\{\mathbf{b}_1, \dots, \mathbf{b}_{n-1}\}$. The algorithm begins by computing the translate $\mathbf{y} \in \Lambda$ of U closest to \mathbf{w} , i.e. the *nearest plane* to \mathbf{w} . The algorithm then implicitly computes the projection \mathbf{w}' of \mathbf{w} onto this translate $U + \mathbf{y}$. It then translates this projection back to the origin (i.e. $\mathbf{w}_{n-1} := \mathbf{w}' - \mathbf{y}$), and solves the smaller problem within U , i.e. the CVP problem $(\Lambda_{n-1}, \mathbf{w}_{n-1})$, where Λ_{n-1} is the lattice formed by the first $n - 1$ basis vectors. The process continues recursively until the final CVP problem with the lattice spanned by \mathbf{b}_1 is solved. The final output is then the sum of the translates from each level of the recursion.

3.4 Problem Statement and Related Work

We have seen that unknown optical path differences (OPD) amongst the apertures in an array, arising from atmospheric turbulence as well as non-idealities in the interferometric system, present a fundamental challenge in interferometry. Moreover, we have seen that while the bispectrum observable eliminates the OPD in interferometric measurements, the ill-posed nature of Fourier phase recovery arising from the OPD terms persists. In Chapter 2 we developed an algorithmic interface which regularizes the ill-posed problem of recovering Fourier phase from the bispectrum by enforcing prior constraints on the image. An alternative approach to prior-regularized reconstruction is to use baseline redundancy to explicitly solve for OPD variation; an array with baseline redundancy contains repeated instances of the same vector baseline involving distinct aperture pairs. Since Fourier phases can be assumed to be equal for all repeated vector baselines, an observed difference amongst their corresponding measurements exposes the contribution of the OPD. This idea of using redundant arrays to calibrate out OPD variation, known as *redundant spacing calibration* (RSC), was developed in works such as those by Arnot *et al.* (1985) and Greenaway (1990). In recent years, innovation in optical technology has engendered a revival of interest in the RSC technique. The simultaneous (or *Fizeau*-style) measurement of fringes on a common

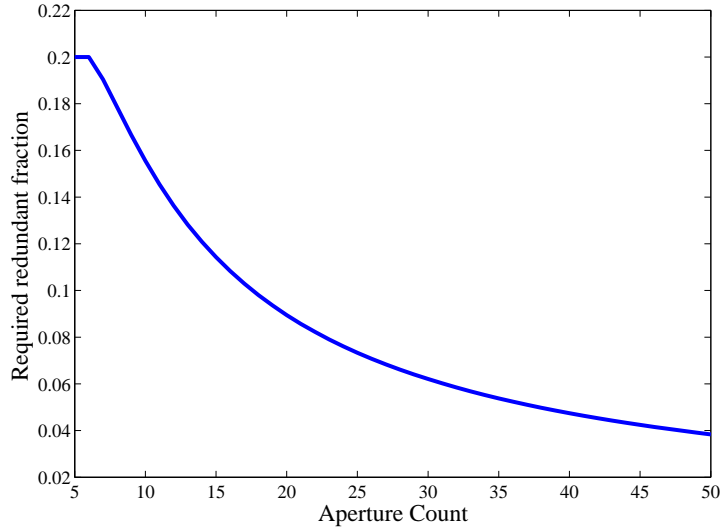


Figure 3.1: Fraction of redundant baselines required for critical redundancy vs. aperture count

focal plane has long been a popular method of acquiring many baseline measurements in an economical manner. However, the Fizeau method had been incompatible with RSC techniques since the fringes formed by each set redundant baselines would alias on the focal plane. An elegant solution to this problem was proposed by Perrin *et al.* (2006). This work developed the idea of segmenting the entrance pupil of a single telescope into an RSC arrangement of sub-pupils from which the light was then coupled via single mode fiber to a non-redundant exit pupil, thereby permitting unambiguous and simultaneous fringe detection for an RSC array. A reconstruction algorithm for this architecture was then proposed in Lacour *et al.* (2007). Even more recently, RSC has been implemented as the calibration scheme of choice for several radio interferometers: the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER) in South Africa (see Ali *et al.* (2015)), the MIT Epoch of Reionization (MITeOR) in the United States (see Zheng *et al.* (2014)), and the Ooty Radio Telescope (ORT) in India (see Marthi and Chengalur (2014)).

As will be shown below, $N - 3$ independent redundant relations are required for unique determination of atmosphere and Fourier phases - a condition we will refer to as *critical redundancy*. An oft-cited drawback of the RSC approach is that it reduces the number of

unique spatial frequencies measured by the interferometer. However, as Figure 3.1 illustrates, the fraction of distinct uv-samples sacrificed for critical redundancy becomes increasingly negligible as the number of apertures in the array increases. Nevertheless the RSC technique presents other challenges which must be overcome for reliable imaging. Central among these challenges is the problem of integer phase ambiguities which arise from the fact that the interferometric phase is only known modulo 2π . Indeed these ambiguities have been shown to play an important role in accurately recovering sensor complex gains and object complex visibilities while imaging with real interferometric instruments (see e.g. Liu *et al.* (2014), Eastwood *et al.* (2009)). In this chapter, we describe these ambiguities and how they can be mitigated using a combination of lattice theory algorithms and careful array design. We will see that these ambiguities have a fundamental presence; namely, they exist whether the calibration strategy works with complex visibilities (which we call the *Phasor* approach) directly, or their respective logarithms (which we call the *Phase* approach). To the best of our knowledge, the results in this chapter are the first to provide array conditions allowing unambiguous interferometric phase determination in spite of wrap ambiguities in the Phase approach, and corresponding false minima in the objective of the Phasor approach.

To motivate the analysis in the chapter, we provide an example illustrating the effect that wrap ambiguities can have in RSC-based image reconstruction. Consider the pattern depicted in the Figure 3.2. This pattern belongs to one of the more popular array classes in the interferometry literature: the so-called *Y-patterns* (see e.g. Arnot *et al.* (1985), Blanchard *et al.* (1996), Labeyrie *et al.* (2006), Eastwood *et al.* (2009), Liu *et al.* (2014)). The corresponding spatial, or *UV*, sampling is provided in the right panel of the Figure.

To demonstrate the potential effect of wrap ambiguities on reconstruction, we simulated both Phase- and Phasor-based reconstruction from noiseless measurements with this pattern. The results are shown in Figure 3.3. The upper left panel shows the true image, and the upper right panel shows the inverse Fourier transform of the UV-sampled visibility function of the object (i.e. the so-called *dirty*, or *interferometric* image). The lower left panel shows the reconstruction result with an implementation of the Phase method similar to that developed

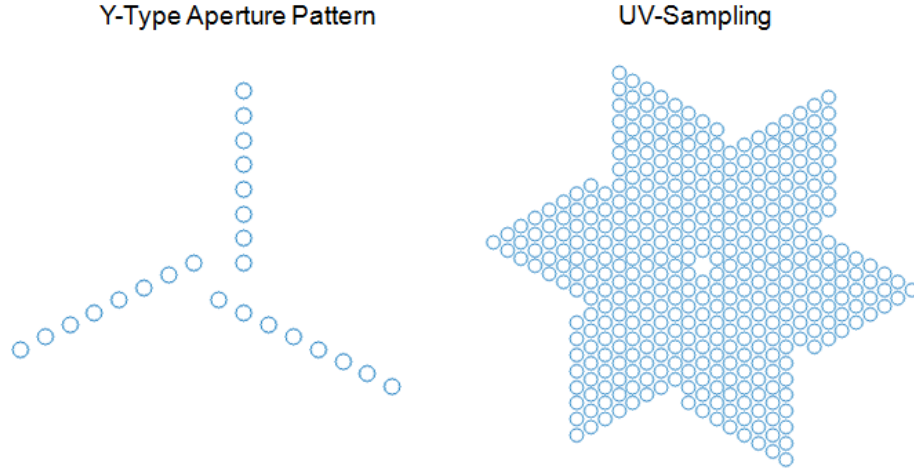


Figure 3.2: *Y-pattern Array Example*

by Lannes (2003), and the lower right panel shows the same for an implementation of the Phasor method (Marthi and Chengalur, 2014) (Lacour *et al.*, 2007). Reconstruction suffers from phase wrapping error in the former case, and a corresponding false-minimum trap in the Phasor case. In the course of this chapter, we will first identify this ambiguity from a mathematical perspective, relate it to a particular physical structure (i.e. the existence of a certain type of loop in the interferometric graph), and provide a simple algorithm for identifying such structures in an arbitrary array so that they can be remedied.

The chapter is organized as follows. In Section 3.5, we review previous work on the integer ambiguity problem, and discuss its presence in the Phase approach. We provide new mathematical conditions for an aperture pattern to be *wrap-invariant*, meaning that the effect of the 2π -periodicity of the measured interferometric phase can be eliminated in image reconstruction. These results are founded upon techniques from lattice theory, as well as the well-known Smith Normal Form (SNF) of an integer matrix. We show the implications of these results on imaging with three types of interferometric observables: the baseline phase measurements, their traditional closure phases, and generalized closure phases. In Section 3.6, we relate these mathematical conditions to conditions on the aperture pattern itself. Namely we show that wrap-invariance is conferred upon arrays satisfying a

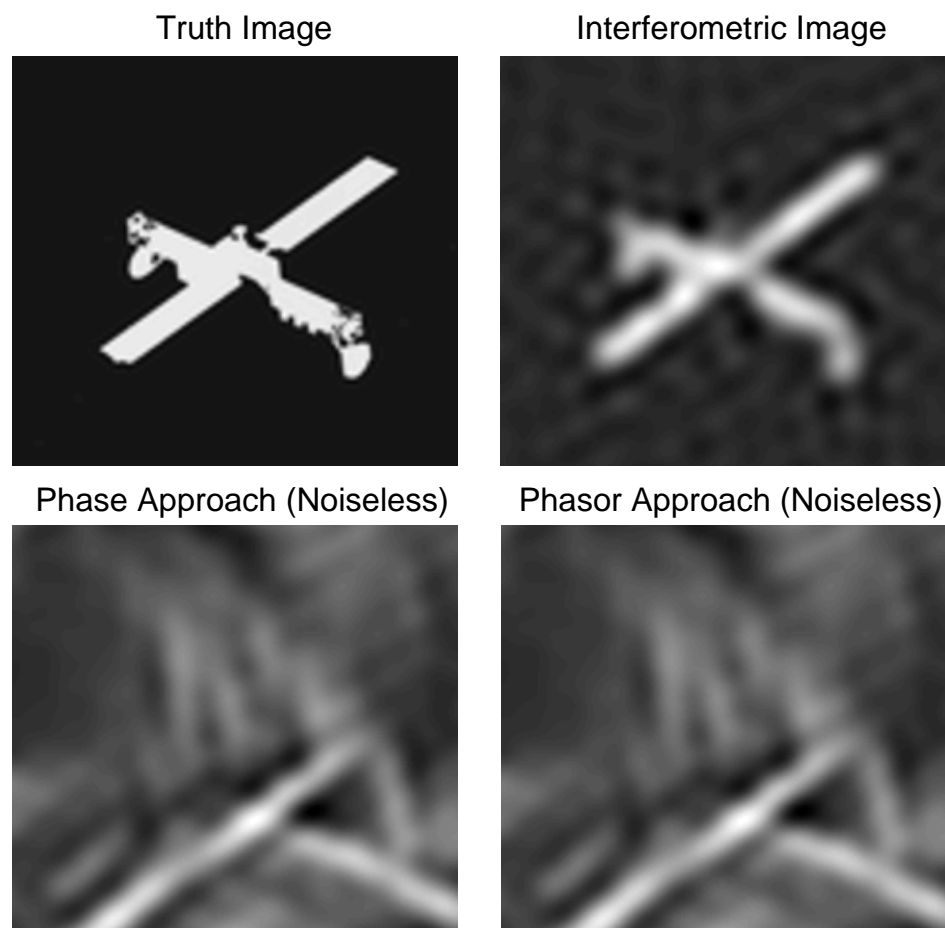


Figure 3.3: *Reconstruction Results for Y-Pattern Example*

certain loop-free condition. As an illustrative example, we diagnose a pattern belonging to the popular Y-pattern class and remedy it to be loop-free. In Section 3.7, we show that the computationally-complex SNF-based approach for ambiguity resolution is actually not necessary for wrap-invariance in many cases, as long as a wrap-induced image shift can be tolerated ¹. For such scenarios, we show that wrap-invariance provides a certificate for the success of existing Phase and Phasor-based approaches in avoiding wrap-induced errors in the former, and false global minima in the latter. Finally we summarize our results in Section 3.8.

3.5 Phase Wrapping Ambiguities in RSC Image Reconstruction

In this Section, we describe a Phase Approach algorithm leveraging the CVP-approach for phase-wrap resolution, which to the best of our knowledge was first developed in Lannes and Anterrieu (1999). We begin by identifying the fundamental phase ambiguity, and subsequently assess its impact on the phase error in the RSC phase solution. In the process, we develop mathematical conditions for wrap-invariance which will form the basis for the notion of a wrap-invariant pattern in Section 3.7. Finally, we assess our results in the context of similar results for approaches requiring the use of closure phases (Lannes, 2003).

3.5.1 Identifying the Fundamental Phase Ambiguity

The traditional approach to RSC reconstruction operates on the measured baseline phases (see e.g. Arnot *et al.* (1985), Greenaway (1994)). To illustrate the approach, let us consider an interferometer which operates at a wavelength λ with two apertures (say, i and j) separated by a vector baseline distance of \mathbf{b}_{ij} . In the absence of any optical path difference, the interference pattern formed by these two apertures encodes a sample of the object's Fourier Transform at spatial frequency $\frac{\mathbf{b}_{ij}}{\lambda}$. Let the true Fourier phase (which we will refer to as

¹As will be shown, this image shift is distinct from the fundamental degeneracy of object position in interferometric measurements

object phase), measured by this interference pattern be denoted as θ_{ij} . The measured phase is given by:

$$\beta_{ij} = \theta_{ij} + \phi_j - \phi_i + 2\pi e \quad (3.2)$$

where $\phi_j - \phi_i$ is the optical path difference between apertures j and i , and e is unknown phase wrap integer arising from the fact that interferometric phase measurements are only known modulo 2π .

Consider an interferometric array which simultaneously makes many such measurements amongst its N apertures. Suppose that of the array's $\binom{N}{2}$ baselines, d of them are distinct. Further suppose we have a solution set $\{\phi_i\}$ and $\{\theta_{ij}\}$ for these equations. Let \mathbf{r}_i denote the vector position of the i -th aperture. As noted by several authors (see, e.g. Wieringa (1992)), we can obtain another valid solution set simply by replacing each ϕ_i with $\phi_i^p = \phi_i + \phi_0 + \mathbf{z} \cdot \mathbf{r}_i$, and each θ_{ij} with $\theta_{ij}^p = \theta_{ij} - \mathbf{z} \cdot (\mathbf{r}_j - \mathbf{r}_i)$, for arbitrary ϕ_0 and \mathbf{z} . Since the free vector \mathbf{z} is a two-parameter vector representing the inherently-ambiguous position of the image within the Field-of-View and the free parameter ϕ_0 is simply a scalar piston offset, the kernel of the RSC system is three-dimensional. This is the *tilt-position degeneracy* which is fundamental in interferometric reconstruction. The end result is that the RSC system contains d unknown distinct object phases and N unknown aperture pistons, and is rank-deficient by at least 3. This implies that there are at most $(d + N - 3)$ linearly-independent equations in the RSC system, and hence at most $N - 3$ redundant relations can be linearly-independent. Remaining relations can be expressed as linear combinations of the elements in this set. A commonly-used simple example of such dependencies (Greenaway, 1994) is shown in Figure 3.4. The phase measurement β_{34} associated with the baseline b_{34} can be expressed as a linear combination of those associated with the other baselines, i.e. $\beta_{34} = \beta_{12} - \beta_{13} + \beta_{24}$.

We will assume for the remainder of the chapter that our array contains $N - 3$ independent relations. Under this assumption, we could in principle solve for a particular solution of this system by arbitrarily setting two object phases (whose spatial frequencies are not co-linear) and one piston phase. This particular solution would then differ from the true

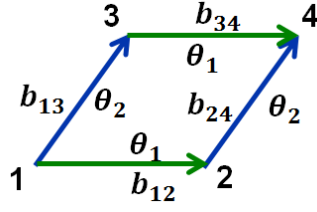


Figure 3.4: Example: *Dependent Parallelogram redundancy.*

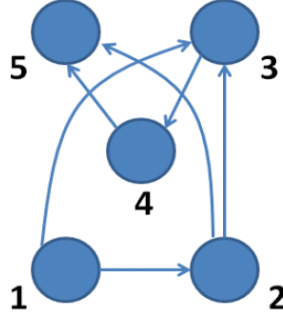


Figure 3.5: Example: *5-aperture RSC pattern. The six distinct baselines are shown.*

solution by a phase ramp in the Fourier domain, corresponding to an image shift in the spatial domain. In this section we will instead construct a different particular solution which is immune to the effects of phase-wrapping by design.

As an example, consider the simple array in Figure 3.5. There are $\binom{5}{2} = 10$ baselines, of which 4 are redundant. A critical array of 5 apertures would have 2 redundancies. Therefore this array possesses more redundancies than necessary (call it *strongly redundant*), and we anticipate that the resulting system will be overdetermined.

The measurement equations associated with this array can be written in matrix form:

$$\begin{bmatrix}
1 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & -1 \\
0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & -1 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & -1 & 0 \\
0 & 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & -1 & 0 \\
-1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & -1 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & -1 \\
0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1
\end{bmatrix}
\begin{bmatrix}
\theta_{12} \\
\theta_{23} \\
\theta_{34} \\
\theta_{45} \\
\theta_{13} \\
\theta_{25} \\
\phi_1 \\
\phi_2 \\
\phi_3 \\
\phi_4 \\
\phi_5
\end{bmatrix}
= \boldsymbol{\beta} + 2\pi\mathbf{e} \quad (3.3)$$

Denoting the matrix above as \mathbf{M} , we can write such systems in compact form as:

$$\mathbf{M} \begin{bmatrix} \boldsymbol{\theta} \\ \boldsymbol{\phi} \end{bmatrix} = \boldsymbol{\beta} + 2\pi\mathbf{e} \quad (3.4)$$

The example above illustrates that the general phase measurement matrix will have two sets of columns: one corresponding to the object phases, and one corresponding to the path differences. We can divide the range of the matrix into two respective subspaces as follows:

Definition 3.5.1. Two fundamental subspaces (Lannes and Anterrieu, 1999): The spectral phase space K is the span of the d columns associated with the Fourier phases, or equivalently the range of the sub-matrix \mathbf{M}_θ formed by these columns. The piston, or *aberration*, phase space L is the span of the N_{ap} columns associated with the piston phases, or equivalently the range of the sub-matrix \mathbf{M}_ϕ formed by these columns.

If we let $n = \binom{N}{2}$ be the number of baselines in the array, the phase measurement matrix \mathbf{M} will be of size $n \times (d + N)$. For a strongly-redundant array like the one in the example above, the column-space $K + L$ of the matrix will clearly not span \mathbb{R}^n . Therefore the wrapped measurement vector $\boldsymbol{\beta}$ will not in general fall in the the subspace $K + L$ (and potentially

can be quite far from it). In the absence of measurement noise, we can unwrap these measurements by identifying those integer correction vectors \mathbf{e} for which $\boldsymbol{\beta}^* = \boldsymbol{\beta} + 2\pi\mathbf{e}$ lies in $K + L$. In the presence of noise, on the other hand, the unwrapped vector will not generally lie in $K + L$ (but for low-to-moderate noise will be in the vicinity). Thus we search for vector(s) $\boldsymbol{\beta}^*$ which are as close to $K + L$ as possible in a weighted least-squares sense (Lannes and Anterrieu, 1999), i.e. we search for the vector

$$\tau_{RSC} = \begin{bmatrix} \hat{\boldsymbol{\theta}}_{RSC} \\ \hat{\boldsymbol{\phi}}_{RSC} \end{bmatrix} = \underset{\mathbf{e}, \boldsymbol{\theta}, \boldsymbol{\phi}}{\operatorname{argmin}} \left\| \mathbf{W} \left(\boldsymbol{\beta}^*(\mathbf{e}) - \mathbf{M} \begin{bmatrix} \boldsymbol{\theta} \\ \boldsymbol{\phi} \end{bmatrix} \right) \right\|^2 \quad (3.5)$$

where \mathbf{W} is the weighting matrix. If we let $\boldsymbol{\Sigma}$ denote the phase measurement covariance matrix and set $\mathbf{W} = \boldsymbol{\Sigma}^{-1}$, this is equivalent to searching for vectors \mathbf{e} which minimize the projection of a whitened measurement $\mathbf{W}^{\frac{1}{2}}\boldsymbol{\beta}^* = \mathbf{W}^{\frac{1}{2}}(\boldsymbol{\beta} + 2\pi\mathbf{e})$ onto the space $(K + L)_{\mathbf{W}}^{\perp} := \ker((\mathbf{W}^{\frac{1}{2}}\mathbf{M})^T)$, where $\ker(\mathbf{A})$ denotes the kernel, or nullspace, of the matrix \mathbf{A} .

Specifically we seek to minimize:

$$f(\mathbf{e}) = \|\mathbf{P}_{\mathbf{W}}\mathbf{W}^{\frac{1}{2}}(\boldsymbol{\beta} + 2\pi\mathbf{e})\|^2 \quad (3.6)$$

where $\mathbf{P}_{\mathbf{W}}$ is a matrix representing the orthogonal projection from \mathbb{R}^n onto $(K + L)_{\mathbf{W}}^{\perp}$.

Letting $\mathbf{e}' = -\mathbf{e}$, we can rewrite the above objective function as:

$$f(\mathbf{e}') = \|\mathbf{P}_{\mathbf{W}}\mathbf{W}^{\frac{1}{2}}(\boldsymbol{\beta} - 2\pi\mathbf{e}')\|^2 = \|\mathbf{P}_{\mathbf{W}}\mathbf{W}^{\frac{1}{2}}\boldsymbol{\beta} - 2\pi\mathbf{P}_{\mathbf{W}}\mathbf{W}^{\frac{1}{2}}\mathbf{e}'\|^2 \quad (3.7)$$

Lannes and Anterrieu (1999) showed that this optimization problem is equivalent to the so-called *closest vector problem* (CVP) in the theory of lattices. We will define a lattice $\mathbf{L}(\mathbb{Z}^n)$ as the set of points generated by integer combinations of the column vectors of a matrix \mathbf{L} . Letting $\tilde{\mathbf{P}} = \mathbf{P}_{\mathbf{W}}\mathbf{W}^{\frac{1}{2}}$, our optimization problem then amounts to the following: Find the lattice point in $\tilde{\mathbf{P}}(\mathbb{Z}^n)$ which is closest to $\tilde{\mathbf{P}}\boldsymbol{\beta}$. Indeed this CVP formulation applies in a very general sense to phase-unwrapping problems taken collectively as discussed in Wubben *et al.* (2011).

A compact representation of the lattice Γ is given by:

$$\Gamma = \left\{ \sum_{i=1}^{m \leq n-(d+N-3)} a_i \mathbf{v}_i \mid \forall a_i \in \mathbb{Z} \right\} \quad (3.8)$$

where $\{\mathbf{v}_i\}$ are linearly-independent and together form a *basis* of the lattice.

Suppose we have found a basis for the lattice $\tilde{\mathbf{P}}(\mathbb{Z}^n)$, and we have solved the Closest Vector Problem for a given measurement vector $\boldsymbol{\beta}$ (c.f. Section 3.3.2). Let \mathbf{b}^* be the output of the Babai Nearest Plane Algorithm - i.e. it is the lattice point which is the closest to $\boldsymbol{\beta}$. We now seek to solve for the wrap vector corresponding to this lattice point, i.e. we seek a solution to:

$$\mathbf{b}^* = \tilde{\mathbf{P}} \hat{\mathbf{e}} \quad (3.9)$$

Note that $\tilde{\mathbf{P}}$ is a (weighted) projection matrix and thus not full-rank, and therefore there will be infinitely-many solutions to this equation. The Closest-Vector-Problem algorithm will provide one particular solution \mathbf{e}_p . The complete set of solutions is then given by:

$$\hat{\mathbf{e}} = \mathbf{e}_p + \mathbf{e}_h \quad (3.10)$$

where \mathbf{e}_h is any integer vector in the kernel of $\tilde{\mathbf{P}}$. Suppose we choose one such vector \mathbf{e}_h and correct our phase measurement vector accordingly. The corrected phase measurement vector can be written as:

$$\hat{\boldsymbol{\beta}}^* = \boldsymbol{\beta} + 2\pi(\mathbf{e}_p + \mathbf{e}_h) \quad (3.11)$$

Lemma 3.5.1. $\mathbf{e}_h \in K + L, \forall \mathbf{e}_h$

Proof: The fact that $\mathbf{e}_h \in \ker(\mathbf{P}_W \mathbf{W}^{\frac{1}{2}})$ implies that $\mathbf{W}^{\frac{1}{2}} \mathbf{e}_h \in \ker(\mathbf{P}_W)$. This in turn implies that $\mathbf{W}^{\frac{1}{2}} \mathbf{e}_h \in \text{im}(\mathbf{W}^{\frac{1}{2}} \mathbf{M})$ and hence that $\mathbf{e}_h \in \text{im}(\mathbf{M})$ since $\mathbf{W}^{\frac{1}{2}}$ is invertible by construction. \square

While any choice of $\mathbf{e}_h \in K + L$ will admit a solution to Equation (3.5), let us consider the optimal one $\mathbf{e}_{h,0}$ which minimizes the error in the ultimate phase solution τ_{RSC} . Let $\boldsymbol{\beta}_0^* = \boldsymbol{\beta} + 2\pi(\mathbf{e}_p + \mathbf{e}_{h,0})$ be the corresponding, optimal unwrapped measurement vector. From Lemma 3.5.1, we see that the unwrapped vector $\hat{\boldsymbol{\beta}}^*$ differs by some $2\pi \mathbf{e}_h^*$ in $K + L$

from β_0^* , i.e.

$$\hat{\beta}^* = \beta_0^* + 2\pi\mathbf{e}_h^* \quad (3.12)$$

The impact of this fundamental ambiguity is the main subject of this chapter. We have depicted the situation in Figure 3.6 to provide a visual summary of the linear algebra involved. As we have seen, the RSC process begins with the interferometric phase measurement β , which due to wrapping will in general lie far from the range $K + L$ of the measurement matrix. By solving the Closest-Vector-Problem using a lattice algorithm such as the Babai Nearest Plane algorithm, it is possible to find a particular correction vector $2\pi\mathbf{e}_p$ which when added to β minimizes the residual in Equation (3.5), i.e. the weighted distance from $K + L$. In the noiseless case, this unwrapped vector $\hat{\beta}^*$ will lie in $K + L$ (i.e. zero residual), whereas in the noisy case, it will in general remain outside of $K + L$.² In either case, the choice of the residual-minimizing vector $\hat{\beta}^*$ is not unique. To see this, let the smallest possible residual norm among all unwrapped candidates be denoted as r_{min} . The set of unwrapped measurement vectors r_{min} away from $K + L$ can be represented as discrete points in a plane parallel to $K + L$, each of which corresponds to a distinct choice of \mathbf{e}_p . Within this family, consider the optimum vector β_0^* whose corresponding least-squares solution τ_{RSC} has minimal error. Any choice of an unwrapped vector in Equation (3.11) is within an error vector $2\pi\mathbf{e}_h^*$ of β_0^* , where \mathbf{e}_h^* is an integer vector in $K + L$. RSC algorithms are fundamentally blind to such errors; distinct unwrappings $\hat{\beta}^*$ and β_0^* both produce solutions to Equation (3.5) in the noiseless case, as do their respective projections onto $K + L$, $\hat{\beta}_{K+L}^*$ and $\hat{\beta}_{0,K+L}^*$, in the noisy case. The property of wrap-invariance introduced in this chapter ensures that the effect of such an error on the phase of the resulting RSC solution τ_{RSC} is either: merely a benign integer multiple of 2π (c.f. Section 3.5.2), or a linear gradient in the estimated Fourier phases, which is equivalent up to an extra image shift (c.f. Section 3.7) in the reconstructed image. In order to develop conditions for wrap-invariance as they relate to pattern design, we must first characterize the effect of the residual wrap error on the RSC

²Without loss of generality, we have selected $\mathbf{e}_h = 0$ in Equation (3.10) so as to simplify the diagram

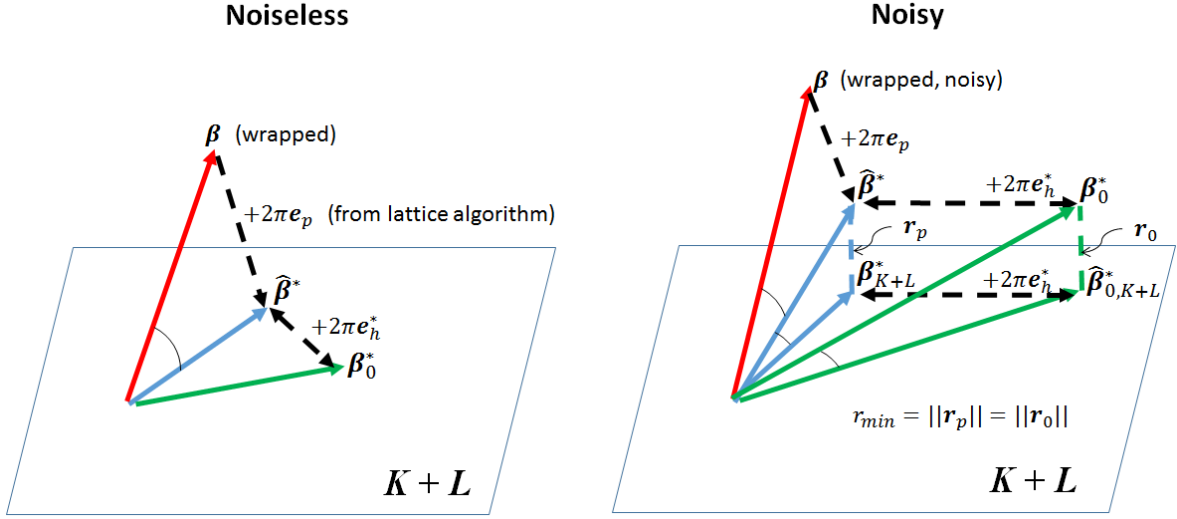


Figure 3.6: Illustration of the fundamental ambiguity of 2π -periodicity in RSC imaging. Distinct unwrappings $\hat{\beta}^*$ and β_0^* both produce solutions to Equation (3.5) in the noiseless case, as do their respective projections onto $K + L$, $\hat{\beta}_{K+L}^*$ and $\hat{\beta}_{0,K+L}^*$, in the noisy case.

least-squares solution. This is the aim of the next section.

3.5.2 Quantifying the Effect of the Fundamental Ambiguity

Let us now quantify the effect of this unresolved wrap error on the ultimate least-squares solution, which can be accomplished in two easy steps. Following standard least-squares principles, we first find the projection $\hat{\beta}_{K+L}^*$ of the unwrapped $\hat{\beta}^*$ onto $K + L$ whose weighted distance from $\hat{\beta}^*$ is minimized. Note that the \mathbf{e}_h^* term in $\hat{\beta}^*$ is already in $K + L$ and hence unchanged by projection. Hence we obtain:

$$\hat{\beta}_{K+L}^* = \mathbf{W}^{-\frac{1}{2}}(\mathbf{I} - \mathbf{P}_W)\mathbf{W}^{\frac{1}{2}}\hat{\beta}^* = \beta_{0,K+L}^* + 2\pi\mathbf{e}_h^* \quad (3.13)$$

where $\beta_{0,K+L}^*$ is the projection of $\hat{\beta}^*$ onto $K + L$. We then solve the system:

$$\mathbf{M} \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \hat{\beta}_{K+L}^* \quad (3.14)$$

As aforementioned, \mathbf{M} is rank-deficient (by 3), and hence there will be infinitely-many

solutions to this system. Successful recovery of the Fourier phases modulo 2π requires a solution preserving the integrality of the error term \mathbf{e}_h^* . In this case, we obtain a final RSC solution $2\pi\mathbf{c}$ away from the true solution for some integer vector \mathbf{c} . To achieve this, we rely on the integer matrix decomposition known as the *Smith Normal Form*, which is described in the following Theorem and Lemma:

Theorem 3.5.2. (*Smith Normal Form*) (Smith, 1861): Let \mathbf{A} be a nonzero $m \times n$ integer matrix with rank r . There exist integer and unimodular³ (and thus invertible) matrices $m \times m$ and $n \times n$ matrices \mathbf{U} and \mathbf{V} respectively such that the matrix product $\mathbf{D} = \mathbf{U}\mathbf{A}\mathbf{V}$ is a diagonal matrix whose diagonal entries \mathbf{D}_{ii} (the so-called elementary divisors) are non-zero integers for $i \leq r$, and zero for $i > r$. Moreover, the matrices \mathbf{U} {and respectively, the matrix \mathbf{V} } represent the row {and column} operations which zeroize \mathbf{A} below {and above} the diagonal.

Theorem 3.5.3. (*Elementary Divisors*) (Smith, 1861): The product of the elementary divisors is the greatest common divisor (gcd) of all $r \times r$ minors of \mathbf{A} .

The proof of the Theorems above can be found in textbooks such as Newman (1972). \square

Let us compute the Smith Normal Form (SNF) $\{\mathbf{U}, \mathbf{D}, \mathbf{V}\}$ of our matrix \mathbf{M} . Let $\mathbf{U}_M = \mathbf{U}^{-1}$ and $\mathbf{V}_M = \mathbf{V}^{-1}$ so that we can write:

$$\mathbf{M} = \mathbf{U}_M \mathbf{D}_M \mathbf{V}_M \quad (3.15)$$

where the r diagonal elements $\{d_i\}$ of \mathbf{D}_M are the elementary divisors of \mathbf{M} .

We can now re-write Equation (3.14) above as:

$$\mathbf{D}_M \mathbf{V}_M \begin{bmatrix} \boldsymbol{\theta} \\ \boldsymbol{\phi} \end{bmatrix} = \mathbf{U}_M^{-1} \hat{\boldsymbol{\beta}}_{K+L}^* \quad (3.16)$$

Let us choose the following particular solution to Equation (3.16):

$$\tau_{RSC} = \mathbf{V}_M^{-1} \mathbf{D}_M^+ \mathbf{U}_M^{-1} \hat{\boldsymbol{\beta}}_{K+L}^* \quad (3.17)$$

³A unimodular matrix is one whose determinant is ± 1

where \mathbf{D}_M^+ denotes the pseudo-inverse of \mathbf{D} . The resulting error is then:

$$\mathbf{e}_{RSC} = \mathbf{V}_M^{-1} \mathbf{D}_M^+ \mathbf{U}_M^{-1} \mathbf{e}_h^* \quad (3.18)$$

Lemma 3.5.4. *Let $\mathbf{u} = \mathbf{U}_M^{-1} \mathbf{e}_h^*$. The residual wrap error \mathbf{e}_{RSC} equals $\vec{0} \bmod 2\pi$ if and only if $\bmod(u_i, d_i) = 0, \forall i \leq r$. The proof of this Lemma is an adaptation of a standard proof which can be found in most textbooks covering linear Diophantine equations (see, e.g. Newman (1972)).*

□

From this Lemma, the following Corollary is clear:

Corollary 3.5.5. *(Sufficient condition on SNF of RSC matrix for wrap-invariance): If the elementary divisors of the measurement matrix \mathbf{M} corresponding to a certain aperture pattern are all 1, the RSC solution defined by τ_{RSC} is immune to phase-wrapping error.*

□

RSC patterns consisting of apertures placed randomly on a Cartesian grid appear to satisfy this sufficient condition with high probability. We conducted a simple experiment in which 15 apertures were randomly placed on a 10×10 grid. Out of 256 placements, 66 were valid RSC patterns (i.e. possessed at least critical redundancy), and of these, only 2 had non-unity elementary divisors.

3.5.3 Relation to closure-phase approaches

The SNF has been applied to the RSC phase problem before (Lannes, 2003). Whereas we have chosen to apply SNF directly to the baseline measurement matrix, the approach taken by Lannes (2003) is to instead treat the piston-invariant phases of the bispectra (the so-called *closure phases*) as the fundamental observables from which the object phases can be inferred via the relation:

$$\mathbf{C}_{o \rightarrow c} \vec{\theta} = \mathbf{y}_{cl} + 2\pi \mathbf{e}_{cl} \quad (3.19)$$

where \mathbf{y}_{cl} are the wrapped closure phases, \mathbf{e}_{cl} is the corresponding wrap vector, and $\mathbf{C}_{o \rightarrow c}$ is the matrix mapping the distinct object phases in the array to closure phases.

Using closure phases as observables can be advantageous in low-light scenarios in which there is not sufficient SNR in a single atmospheric coherence time to reliably measure the baseline phases. To overcome this low per-frame SNR, atmosphere-invariant observables such as the bispectra can be integrated over many frames to build sufficient SNR, and their respective closure phases used as reliable phase measurements. Since the baseline phases are known modulo 2π , the linear combinations of them which comprise the closure phases are also known modulo 2π .

Lannes (2003) applies the SNF to the closure matrix $\mathbf{C}_{o \rightarrow c}$. By direct analogy to Corollary 3.5.5, note that if the elementary divisors of \mathbf{C}_{oc} are all 1, then the RSC solution is immune to phase-wrapping error. Otherwise, severe distortion is possible in the resulting image reconstruction. In order to relate this condition to Corollary 3.5.5, let us first define another closure matrix $\mathbf{C}_{m \rightarrow c}$ which instead maps the phase measurements to closure phases. This mapping consists of equations of the form:

$$y_{123} = \beta_{12} + \beta_{23} - \beta_{13} \quad (3.20)$$

where y_{123} is the closure phase associated with apertures 1, 2, and 3, and the β_{ij} are the associated baseline phases (see Equation (3.2)). Of the $\binom{N}{3}$ possible closure phases, at most $\binom{N-1}{2}$ can be linearly-independent (see e.g. Readhead *et al.* (1988)). One commonly-chosen set of such linearly-independent relations consists of all the closure triangles involving a given aperture A , and this is the set selected by Lannes (2003). $\mathbf{C}_{m \rightarrow c}$ is therefore an $\binom{N-1}{2} \times \binom{N}{2}$ matrix. Lannes (2003) accordingly provides a convenient grouping of the baselines into two categories: (1) *spanning tree* baselines which connect aperture A to all other apertures, and (2) *loop entry* baselines which provide the closure for these spanning tree baselines. This categorization is depicted in Figure 3.7.

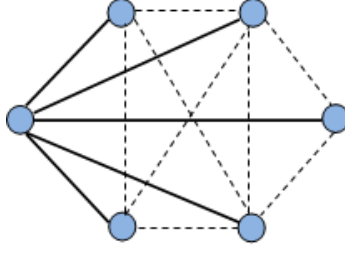


Figure 3.7: Distinction between spanning tree baselines (thick, solid) and loop entry baselines (thin, dotted)

Given this categorization, we can decompose $\mathbf{C}_{m \rightarrow c}$ into corresponding blocks as:

$$\mathbf{C}_{m \rightarrow c} = \begin{bmatrix} \hat{\mathbf{C}}_{m \rightarrow c} & \mathbf{I}_{\binom{N-1}{2}} \end{bmatrix} \quad (3.21)$$

where $\hat{\mathbf{C}}_{m \rightarrow c}$ contains the spanning tree contributions to the matrix (which appear in multiple closures), and $\mathbf{I}_{\binom{N-1}{2}}$ is the $\binom{N-1}{2} \times \binom{N-1}{2}$ identity matrix representing the loop-entry contributions (each of which appears in only one closure). The following property follows from this block form expression:

Lemma 3.5.6. *The elementary divisors of $\mathbf{C}_{m \rightarrow c}$ are all 1.*

Proof: Since we have chosen a linearly-independent subset of closure relations, $r = \text{rank}(\mathbf{C}_{m \rightarrow c}) = \binom{N-1}{2}$. There exists a $r \times r$ minor (namely, $\mathbf{I}_{\binom{N-1}{2}}$) which is equal to 1. Therefore the gcd of all $r \times r$ minors is 1, and therefore from Theorem 3.5.3 (Elementary Divisors), all elementary divisors must be 1. \square

Let us now relate $\mathbf{C}_{m \rightarrow c}$ to the matrix $\mathbf{C}_{o \rightarrow c}$ used by Lannes (2003). Recall from the discussion of bispectra in Section 3.4 that the closure relations eliminate piston differences in the measurements so that $\mathbf{C}_{m \rightarrow c}$ annihilates the subspace L , i.e. the space spanned by the columns of \mathbf{M} corresponding to $\vec{\phi}$. Defining \mathbf{M}_θ as the submatrix of \mathbf{M} containing the columns corresponding to $\vec{\theta}$, we have

$$\mathbf{C}_{m \rightarrow c} \mathbf{M} \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{o \rightarrow c} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \theta \\ \phi \end{bmatrix} \quad (3.22)$$

where $\mathbf{C}_{o \rightarrow c} = \mathbf{C}_{m \rightarrow c} \mathbf{M}_\theta$. $\mathbf{C}_{o \rightarrow c}$ is an $\binom{N-1}{2} \times d$ matrix which is rank-deficient by two ⁴.

Then by direct analogy to Equation (3.14), we can obtain valid RSC object phase solutions by solving:

$$\mathbf{C}_{m \rightarrow c} \mathbf{M} \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{o \rightarrow c} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \mathbf{y}_{cl}^* + 2\pi \mathbf{e}_{h,cl}^* \quad (3.23)$$

where \mathbf{y}_{cl}^* is the true unwrapped closure vector and $2\pi \mathbf{e}_{h,cl}^*$ is the residual integer wrapping error vector after applying the Babai NP algorithm to solve the CVP problem associated with matrix $\mathbf{C}_{m \rightarrow c} \mathbf{M}_\theta$ and $2\pi \mathbf{e}_{cl}$. Note that if we find a vector $\vec{\theta}^*$ satisfying Equation (3.23), it will clearly also satisfy the relation $\mathbf{C}_{o \rightarrow c} \vec{\theta}^* = \mathbf{y}_{cl}^* + 2\pi \mathbf{e}_{h,cl}^*$ of Lannes (2003). Note furthermore that we can solve the equation above in two separate integer-preserving steps of the form of Equation (3.17), the first involving the SNF decomposition of $\mathbf{C}_{m \rightarrow c}$, and the second involving that of \mathbf{M} . Since the elementary divisors of $\mathbf{C}_{m \rightarrow c}$ are all 1 by construction (by Lemma 3.5.6) and hence the first step is thus integer-preserving, wrap-invariance again amounts to whether or not all elementary divisors of \mathbf{M} are 1. Therefore we have the following Proposition relating wrap invariance for closure measurements to that for raw phase measurements:

Proposition 3.5.7. (*Sufficient condition for wrap-invariance of closure-based RSC*): *If the elementary divisors of the phase measurement matrix \mathbf{M} are all 1, then the closure-based RSC solution will be wrap-invariant.*

Proof: (see Appendix B.1) □

We remark in passing that although the preceding analysis was presented in the context of the traditional three-aperture closure, it applies directly to the case of closures involving an arbitrary number of sides. As an example, consider the pattern shown in Figure 3.8. A spanning tree for the pattern consisting of the short baselines in an array is depicted. Let $\{\phi_{sp}\}$ denote the aperture phase differences in these $N - 1$ spanning tree baselines. Note

⁴To see this, note that each solution set to Equation (4.49) above remains valid after replacing each θ_{ij} with $\theta_{ij}^p = \theta_{ij} - \mathbf{z} \cdot (\mathbf{r}_j - \mathbf{r}_i)$

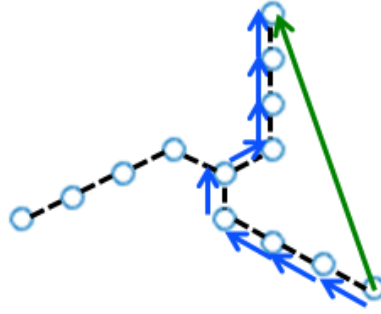


Figure 3.8: Bootstrapping phase of a low-SNR baseline (green) with subset (blue) of high-SNR baselines from spanning tree baselines (black)

that all aperture phase differences in the array can be expressed as linear combinations of the $\{\phi_{sp}\}$. If the aperture phase differences are known reliably via measurements of the spanning tree baselines, we can use these measurements to cancel the aperture phase differences in all other measurements (of which one example is shown in green). The idea of using such *generalized closure phases* (Martinache, 2010) is indeed the mathematical foundation for the promising technique known as *baseline bootstrapping* in which high-fidelity phase measurements of several high-SNR baselines (typically the short baselines) to cancel the atmosphere on each lower-SNR (and hence lower fidelity) baseline.

Note that for an arbitrary N -aperture pattern, there will in general be $N - 1$ spanning tree baselines and thus $\binom{N}{2} - (N - 1) = \binom{N-1}{2}$ generalized closures, each involving a distinct closing (or loop-entry) baseline. Therefore the resulting measurement matrix can be expressed exactly as in Equation (3.21) above and hence the preceding analysis holds.

While this section has considered a few possibilities for phase observables, relating mathematical conditions for wrap invariance to a physical condition on aperture placement is more intuitive when considering the raw phase measurements as opposed to their closures. For this reason, for the remainder of the chapter we will work directly with the baseline phases and their associated wrapping errors. In particular, we will begin by connecting these wrapping errors with analogous ambiguities in recently-developed phasor-based approaches.

3.6 Implications of Wrap Ambiguities on Pattern Design

In this section we use the mathematically-sufficient conditions for wrap invariance from the previous section to show that aperture patterns whose interferometric graph satisfies a certain loop-free condition are wrap-invariant. Here we define the *interferometric graph* in the standard way: it is simply the graph formed by connecting the array's apertures (the *nodes* of the graph) with edges representing the array's baselines.

This condition is founded on the Theorem 3.5.3 (Elementary Divisors) in Section 3.5 and the following definition of the matrix determinant. This definition is given in many linear algebra texts (see e.g. Bretscher (2001)).

Definition 3.6.1. Matrix Determinant in terms of patterns: Suppose we have an $n \times n$ matrix \mathbf{A} . Define a *pattern* as a selection of n entries of the matrix in which there is only one chosen entry in each row and one in each column of the matrix. Furthermore, we denote a pair of numbers in a pattern as *inverted* if one of them is located above and to the right of the other. Then we can obtain the determinant of \mathbf{A} by summing the products associated with all patterns with an even number of inversions and subtracting the products associated with all the patterns with an odd number of inversions.

Consider a $r \times r$ sub-matrix $\tilde{\mathbf{M}}_I$ consisting of a set I of linearly-independent rows in \mathbf{M} and $d + N - 3$ linearly-independent columns. Our goal will be to find conditions under which such a sub-matrix contains only one pattern with a non-zero product, in which case the determinant will be ± 1 from the definition above. This will guarantee, by the Theorem 3.5.3 (Elementary Divisors) of the previous section, that the elementary divisors of \mathbf{M} are all 1, which will in turn ensure that the error in the RSC solution τ_{RSC} will be $0 \pmod{2\pi}$ by Corollary 3.5.5.

We begin our search for such a matrix by forming the intermediate $r \times (r + 3)$ matrix \mathbf{M}_I consisting of I linearly-independent rows in \mathbf{M} , and define the corresponding sub-graph G containing only the baselines corresponding to these rows. In particular, let us examine the pattern restrictions encountered if we select a given column in \mathbf{M}_I for participation in

our $r \times r$ sub-matrix $\tilde{\mathbf{M}}_I$. Namely, we will sequentially identify those special matrix entries which must be part of a pattern with a non-zero product. For example, all non-redundant measurements contain a singleton ± 1 in the column associated with their object phase. All non-zero patterns must clearly contain this ± 1 and so we can select these singleton object-phase entries as guaranteed participants in a non-zero pattern. Moreover, all measurements containing a *leaf node* (i.e. a node with a single connection) in G contain a singleton ± 1 in the column associated with their leaf node. All non-zero patterns must clearly contain this ± 1 as well. Thus we can also select these leaf node entries as guaranteed participants in a non-zero pattern.

There may be cascading implications of such singleton measurements. To illustrate this, consider the scenario shown in Figure 3.9. A simple RSC array is shown on the left. A subset of the baselines in one possible linearly-independent sub-matrix \mathbf{M}_I is depicted. A simplified depiction of the matrix \mathbf{M}_I is shown in which all non-zero entries have been colored black and all zero entries have been colored white for simplicity. In Step A of the reduction process, object phase θ_5 is selected for participation since it is a singleton object phase. Its corresponding row (i.e. row 5) in \mathbf{M}_I is then eliminated from participation since the remaining entries in this row cannot participate in a pattern (by definition of a pattern). In Step B, the aperture 6 entry ϕ_6 in row 9 is selected since it has become a leaf node in the pattern, and row 9 can then be eliminated. Then in Step C, object phase θ_4 is then selected by virtue of becoming a singleton object phase, and row 4 is then eliminated. This selection/elimination process can be repeated beyond the steps shown in the Figure, until either no leaf nodes and singleton object phases remain, or there are no baselines left to eliminate. We formalize the *pattern reduction* process in the listing Algorithm 2.

We can see that any baseline in an interferometric graph which does not belong to a loop will be eliminated in the reduction process. The only structures in the graph that persist after this reduction are sets of loops with a certain property. Namely we define a *persistent loop set* as a set of loops that contains at least two instances of every baseline contained in the set. (A set can consist of any number of loops, including one). With this definition,

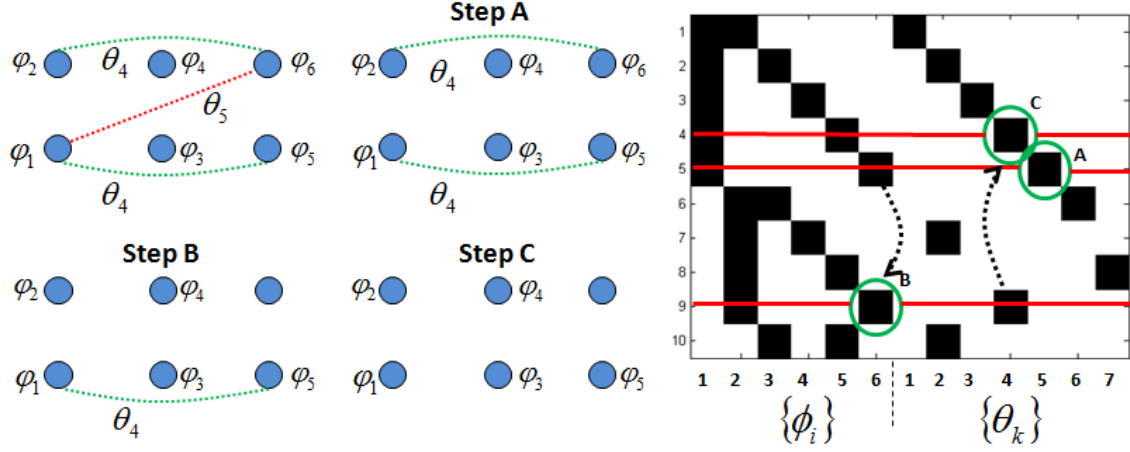


Figure 3.9: Example: Reducing an aperture pattern and associated matrix to identify Persistent Loop(s)

Algorithm 2 Pattern Reduction Algorithm

Require: R {where R is the set of the baselines corresponding to \mathbf{M}_I , where I denotes the indices of a linearly-independent subset of $d + N - 3$ rows of \mathbf{M} }

while $|R| > 0$ **do**

1. **Leaf Nodes**
 - 1.1 remove any remaining baselines containing leaf nodes from R
 - 1.2 add the associated apertures to the list N
2. **Singleton Object Phases**
 - 2.1 remove any remaining baselines containing singleton object phases from R
 - 2.2 add the associated object phases to the list O

if no baselines removed in the current iteration **then**

return PERSISTENT

end if

end while

return LOOPFREE

absolute invariance may be possible if the graph of the redundant baselines does not contain any persistent loop sets. Algorithm 2 returns PERSISTENT if persistent loops exist and LOOPFREE if the pattern is completely reduced and therefore free of persistent loops.

Note that in the latter case, we will have eliminated r rows from \mathbf{M}_I . Since each baseline elimination is associated with object phase or aperture selection from distinct columns, and \mathbf{M}_I contains $r + 3$ columns, there will be exactly three extraneous columns not involved in the reduction process. The $r \times r$ submatrix obtained by selecting the non-extraneous columns (i.e. those corresponding to the selected object phases and leaf nodes in O and N , respectively) will then have a determinant of ± 1 by virtue of having a single non-zero pattern revealed by the reduction process. Having ensured the existence of a unit $r \times r$ minor, the gcd of all $r \times r$ minors must be 1. We have hence confirmed the elementary divisors must be all 1, and thereby ensured that the pattern is wrap-invariant (c.f. Theorem 3.5.3 (Elementary Divisors), and Corollary 3.5.5). We summarize the sufficient condition as follows:

Proposition 3.6.1. (Sufficient conditions on aperture pattern for wrap-invariance): *Consider the graph of an aperture pattern which contains d distinct baselines and any set of $N - 3$ linearly-independent redundant baselines. If this graph does not contain persistent loop sets (in the sense defined above), the matrix $\tilde{\mathbf{M}}_I$ formed by these independent measurements will have determinant ± 1 . As a result Corollary 3.5.5 will hold, thereby guaranteeing that RSC solution τ_{RSC} will be invariant to wrapping of the phase measurements.*

We have hence arrived at a physical definition of a *wrap-invariant* pattern. We now apply Algorithm 2 to the example pattern shown in Figure 3.2. Algorithm 2 reduces the pattern to the persistent loop set shown in Figure 3.10.

The elementary divisors of the pattern's measurement matrix are not all 1; they are all 1 except for a singleton 3 and hence $\det(\tilde{\mathbf{M}}_I) \bmod 3 = 0$ for all choices of $\tilde{\mathbf{M}}_I$.

Having traced the distortion induced by phase wrapping to physical property of the array itself, we now return to Figure 3.3. The lower left panel shows the reconstruction result with the SNF-based Phase method described in Section 3.5. The closure phase approach

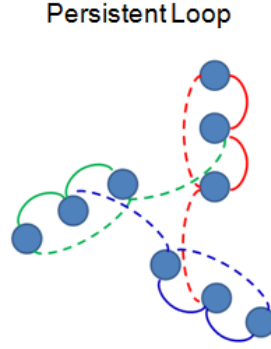


Figure 3.10: *Persistent Loop set at center of pattern in Figure 3.2*

yields the same corruption in reconstruction, as the elementary divisors of \mathbf{C}_{oc} are also all 1 except for a singleton 3.

There are several simple ways to amend this pattern so that it is wrap-invariant. While the most intuitive of these involve moving the apertures involved in the persistent loop shown in Figure 3.10, these approaches leave gaps in the UV-sampling pattern. An alternate approach that preserves the UV-sampling is to add an aperture to the center of the pattern as shown in Figure 3.11. This results in additional linearly-independent redundancies colored in blue and green, respectively, in the Figure. These additions replace baselines in the persistent loop, allowing this loop to be broken. With wrap-invariance, reconstruction results match the true image in both the phase and phasor approaches as respectively shown in Figure 3.12. In the top row, reconstruction results are displayed for the phase (left) and phasor (right) approaches for the noiseless case. The image distortion present in Figure 3.3 has been completely eliminated by tweaking the pattern so that it is wrap-invariant. Analogous results for an SNR of 25 dB are displayed in the bottom row. Here we define SNR as the ratio of the phasor magnitude at visibility 1 (i.e. zero spatial frequency) to the standard deviation of the noise, which we have assumed to be complex Gaussian and i.i.d. across spatial frequency for this simulation.

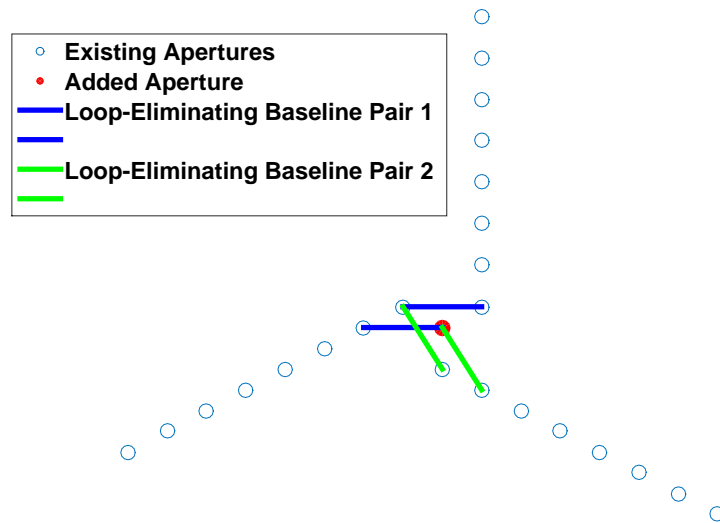


Figure 3.11: *Amended Pattern*

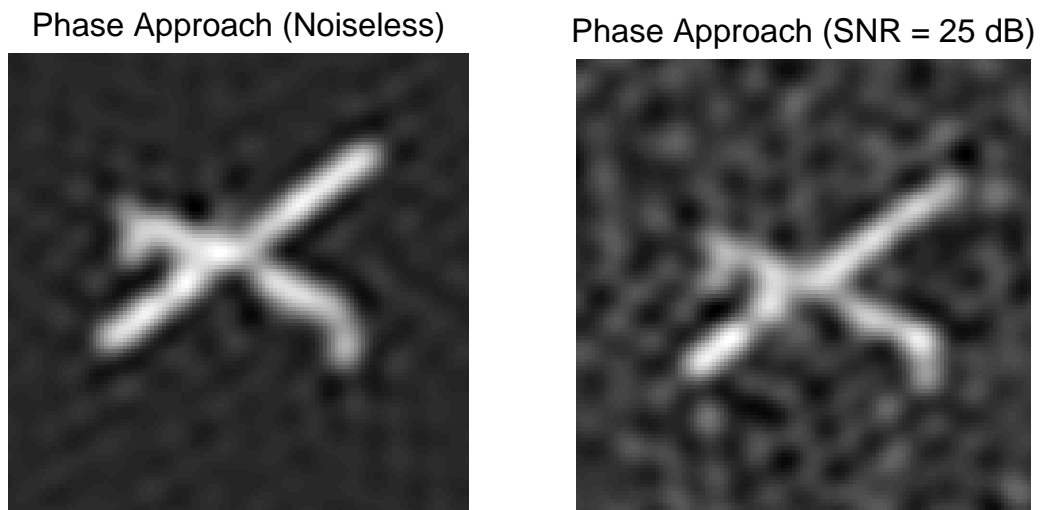


Figure 3.12: *Reconstruction Results for Amended Pattern*

3.7 Wrap-invariance and Practical RSC Calibration

In previous sections we established that the Phase Approach can be made robust to phase-wrapping using the Smith Normal Form (SNF) and algorithms from lattice theory. Moreover, the SNF provided a mathematical framework for the notion of wrap-invariance. From a practical standpoint, however, computation of the Smith Normal Form is likely to become a computational burden for large arrays, as in those under current consideration with $N_{ap} \approx 10^2$ and $n \approx 10^4$ baselines (Zheng *et al.*, 2014). Techniques not requiring such a computation are hence of strong practical interest. In this section we show that the wrap-invariance property checked by Algorithm 2 provides a certificate for reliable reconstruction with such techniques, in the presence of the fundamental 2π -periodicity of interferometric measurements.

3.7.1 Practical Phase Approaches

Our approach here will rely upon the well-known Singular Value Decomposition (SVD) of the measurement matrix \mathbf{M} , which is given by:

$$\mathbf{M} = \mathbf{U}_\sigma \mathbf{\Sigma}_\sigma \mathbf{V}_\sigma^T \quad (3.24)$$

in which \mathbf{U}_σ and \mathbf{V}_σ are $m \times m$ and $(d + N) \times (d + N)$ orthogonal matrices, respectively. $\mathbf{\Sigma}_\sigma$ is a $m \times (d + N)$ diagonal matrix with r non-zero diagonal entries (the so-called *singular values of \mathbf{M}*), where $r = \text{rank}(\mathbf{M}) = d + N - 3$.

Lemma 3.7.1. *The final 3 columns of \mathbf{V}_σ form a basis for the nullspace of \mathbf{M} .*

Proof: This follows from the fact \mathbf{M} is rank-deficient by 3, and standard properties of the right singular vectors comprising \mathbf{V}_σ in the SVD. (Bretscher, 2001) \square

Now recall that in Section 3.4, we provided an one particular, SNF-based, solution to Equation (3.14). Here we instead consider the complete set of solutions to this Equation, which is given by:

$$\boldsymbol{\tau}_\sigma = \mathbf{M}^+(\hat{\boldsymbol{\beta}}_{0,K+L}^* + 2\pi\mathbf{e}_h^*) + \boldsymbol{\tau}_0 \quad (3.25)$$

where $\boldsymbol{\tau}_0$ is any vector in the nullspace of \mathbf{M} , and \mathbf{M}^+ denotes the pseudo-inverse of \mathbf{M} , whose matrix elements are derived directly from the SVD above.

$$\mathbf{M}^+ = \mathbf{V}_\sigma \boldsymbol{\Sigma}_\sigma^+ \mathbf{U}_\sigma^T \quad (3.26)$$

where $\boldsymbol{\Sigma}_\sigma^+$ is a $(d + N) \times m$ diagonal matrix whose r non-zero diagonal entries are the reciprocals of the corresponding non-zero entries in $\boldsymbol{\Sigma}_\sigma$.

Typically Phase Approach techniques implicitly select a particular solution from the family of solutions in Equation (3.25) via augmenting the matrix \mathbf{M} with additional constraints. The most common among these (Wieringa, 1992) (Wijnholds and Noorishad, 2012) is to enforce that: $\sum \phi = 0$, $\sum \phi_i \mathbf{r}_i = \mathbf{0}$, where \mathbf{r}_i is the vector position of the i -th aperture in the array.

Note that the error resulting from application of this pseudo-inverse to the unwrapped measurement vector will be given by:

$$2\pi\mathbf{e}_\sigma = \mathbf{M}^+(2\pi\mathbf{e}_h^*) \quad (3.27)$$

Note furthermore that the vector representing the solution error \mathbf{e}_σ has two parts: one for the error in the Fourier phases (which we denote with the index K), and one for the error in the atmosphere/piston (which we will denote with the index L), i.e.:

$$\mathbf{e}_\sigma = [\mathbf{e}_{\sigma,K}, \mathbf{e}_{\sigma,L}]^T \quad (3.28)$$

We will focus attention on the latter, since it is of direct relevance for image formation. Let us express the spatial frequencies measured by an array as two-element vectors of the form (ω_x, ω_y) . Let \mathbf{X} be the $d \times 2$ matrix containing these spatial frequencies. Note then that the phase-wrap error will manifest itself merely as an image shift if and only if this error is a (modulo- 2π) phase ramp, i.e. there exists a 2-element shift vector \mathbf{z} and an integer vector

\mathbf{k} which together satisfy:

$$2\pi\mathbf{e}_{\sigma,K} - 2\pi\mathbf{X}\mathbf{z} = 2\pi\mathbf{k} \quad (3.29)$$

Substituting from Equation (3.27) we obtain:

$$\mathbf{M}_K^+(2\pi\mathbf{e}) - 2\pi\mathbf{X}\mathbf{z} = 2\pi\mathbf{k} \quad (3.30)$$

where \mathbf{M}_K^+ denotes the sub-matrix of \mathbf{M}^+ formed by the rows associated with K .

Dividing through by 2π we obtain the equation: $\mathbf{M}_K^+\mathbf{e}_h^* - \mathbf{X}\mathbf{z} = \mathbf{k}$. Note that each element of \mathbf{M}_K^+ can be expressed as some rational number $\frac{p_i}{q_i}$. Similarly we first assume \mathbf{X} contains rational spatial frequencies with greatest common denominator q_x . Then we can multiply through by the least-common-multiple (LCM) of the $\{q_i\}$ and q_x to obtain a system of equations whose coefficients are guaranteed to be integer (i.e., we have a linear Diophantine system). Let this LCM be denoted as l . Then we have, after rearranging terms,

$$l\mathbf{X}\mathbf{z} = l(\mathbf{M}_K^+\mathbf{e}_h^* - \mathbf{k}) \quad (3.31)$$

We now wish to determine conditions under which there exist vectors \mathbf{k} and \mathbf{z} satisfying this overdetermined Diophantine system. Applying the Smith Normal Form decomposition (c.f. Theorem 3.5.2) to the matrix $l\mathbf{X}$ this time, and noting that $\text{rank}(\mathbf{X}) = 2$, we have:

$$\mathbf{D}_X = \mathbf{U}_X(l\mathbf{X})\mathbf{V}_X \quad (3.32)$$

where \mathbf{U}_X and \mathbf{V}_X are unimodular matrices of size $d \times d$ and 2×2 , respectively, and \mathbf{D}_X is a rectangular diagonal matrix whose entries are zero below row 2.

If we left-multiply Equation (3.31) by \mathbf{U}_X on both sides, we obtain:

$$l\mathbf{U}_X\mathbf{X}\mathbf{z} = l\mathbf{U}_X(\mathbf{M}_K^+\mathbf{e}_h^* - \mathbf{k}) \quad (3.33)$$

Using Equation (3.32) and the fact that \mathbf{V}_X is a unimodular (and hence invertible) matrix, we can then write:

$$\mathbf{D}_X \mathbf{V}_X^{-1} \mathbf{z} = l(\mathbf{U}_X \mathbf{M}_K^+ \mathbf{e}_h^* - \mathbf{U}_X \mathbf{k}) \quad (3.34)$$

We are now in position to prove the main result of this section, which is preceded by the following Lemma:

Lemma 3.7.2. *Given wrap-invariance, the vector $\mathbf{U}_X \mathbf{M}_K^+ \mathbf{e}_h^*$ has integer entries below row 2.*

Proof: (see Appendix B) □

Proposition 3.7.3. *If a pattern is wrap-invariant (in the sense of Section 3.5), reconstruction error induced by phase wrapping is limited to an image shift.*

Proof: We re-arrange Equation (3.34) above so that it reads:

$$\frac{1}{l} \mathbf{D}_X \mathbf{V}_X^{-1} \mathbf{z} - \mathbf{U}_X \mathbf{M}_K^+ \mathbf{e}_h^* = -\mathbf{U}_X \mathbf{k} \quad (3.35)$$

Let $\mathbf{v} = \frac{1}{l} \mathbf{D}_X \mathbf{V}_X^{-1} \mathbf{z} - \mathbf{U}_X \mathbf{M}_K^+ \mathbf{e}_h^*$. Note that since \mathbf{D}_X is zero below row 2, the entries of \mathbf{v} below row 2 will be equal to those of $(-\mathbf{U}_X \mathbf{M}_K^+ \mathbf{e}_h^*)$, which are integers by Lemma 3.7.2. Now consider the first and second entries of \mathbf{v} . Let \mathbf{f} be the vector containing the fractional parts of the first two elements of vector $\mathbf{U}_X \mathbf{M}_K^+ \mathbf{e}_h^*$, and let \mathbf{A} be the invertible matrix consisting of the first two rows of $\frac{1}{l} \mathbf{D}_X \mathbf{V}_X^{-1}$. Without loss of generality, choose $\mathbf{z}^* = \mathbf{A}^{-1} \mathbf{f}$ so that the fractional part \mathbf{f} is annihilated, leaving only integer elements in the first two entries of \mathbf{v} . Hence we now have:

$$\mathbf{v} = -\mathbf{U}_X \mathbf{k} \quad (3.36)$$

with \mathbf{v} ensured to contain only integer elements. Since \mathbf{U}_X is unimodular, the vector $\mathbf{k}^* = -\mathbf{U}_X^{-1} \mathbf{v}$ will be integral. We have thus found a pair $(\mathbf{z}^*, \mathbf{k}^*)$ with integer \mathbf{k}^* which satisfies the Equation (3.34). Since Equation (3.34) is related to Equation (3.31) via a unimodular (and hence invertible) mapping \mathbf{U}_X , invariance is hence proven. □

With the previous result, we have characterized the complete family of Phase Approach solutions given in Equation (3.25). Namely we have shown that, for a wrap-invariant pattern,

the family differs by at most an image shift from the true solution. Returning to our running example in Figure 3.11, we verified that different solution choices from the family given in Equation (3.25) simply resulted in shifts of an otherwise pristine image in the reconstruction. On the other hand with wrap-variant pattern in Figure 3.2, image distortion of the severity of Figure 3.3 was again observed, as expected.

3.7.2 Practical Phasor Approaches

Though traditional treatments employ the phase approach of the previous section which operates on baseline phases, recent papers (e.g. Marthi and Chengalur (2014), Liu *et al.* (2010)) have shown that approaches which operate at the phasor level can be superior in accuracy. Liu *et al.* (2010) developed a Gauss-Newton-type Non-linear Least-Squares (NLS) solver and showed it produced unbiased phase estimates, in contrast with the biased ones provided by the phase approach. Marthi and Chengalur (2014) and Wijnholds and Noorishad (2012) have also proposed low-complexity phasor-based approaches and demonstrated performance near the Cramer-Rao Bound. Though the capacity of the Phasor approaches to produce superior accuracy relative to Phase approach has been demonstrated, the former's convergence issues can be mitigated via initialization with the results of the latter (see e.g. Liu *et al.* (2010), and Zheng *et al.* (2014)).

The implementations of the Phasor Approach typically employ the following measurement model:

$$V_{ij} = g_i g_j^* f_{ij} + n_{ij} \quad (3.37)$$

where V_{ij} is the complex visibility observed between apertures i and j , $g_i = |g_i|e^{j\phi_i}$ and $g_j = |g_j|e^{j\phi_j}$ are the complex gains of these apertures, f_{ij} is the true complex visibility measured by this pair, and n is complex measurement noise. Note that the phase difference between g_i and g_j is simply the optical path difference between apertures i and j introduced in the previous section. Given this model, NLS approaches attempt to find a set of complex phasors $\{g_i\}$ and $\{f\}$ which minimize an objective function of the form:

$$\Lambda = \sum_i \sum_{j>i} w_{ij} ||(V_{ij} - g_i g_j^* f_{ij})(V_{ij}^* - g_i^* g_j f_{ij}^*)|| \quad (3.38)$$

Minimization of Λ with respect to the unknowns (i.e. distinct object and antenna complex gains) can be accomplished with iterative application of the following updates, as reported by Marthi and Chengalur (2014) in the context of radio interferometry, and by Lacour *et al.* (2007) in the context of optical interferometry:

$$g_k = \frac{\sum_{j \neq k} w_{kj} g_j^* f_{kj}^* V_{kj}}{\sum_{j \neq k} w_{kj} |g_j|^2 |f_{kj}|^2} \quad (3.39)$$

$$f_b = \frac{\sum_{j>k} g_k^* g_j V_{kj}}{\sum_{j>k} w_{kj} |g_k|^2 |g_j|^2} \quad (3.40)$$

where the $\{f_b\}$ are the true complex visibilities of the distinct object phases in the array.

Due to the circularity of these definitions, these equations must be solved iteratively. Starting from an initial guess for all phasors, Equation (3.39) is solved to obtain a better estimate for the $\{g_k\}$ and then these $\{g_k\}$ are used to obtain refined estimates of the $\{f_b\}$ through Equation (3.40). In the next iteration, these $\{f_b\}$ are used to further refine $\{g_k\}$, and so on.

Though there are other means of minimizing objectives of the form Λ (Wijnholds and Noorishad, 2012) (Liu *et al.*, 2010), we omit discussion of them here; our present purpose is to characterize the correctness of the solutions themselves, regardless of how they are obtained. As has been noted before (see e.g. Lannes and Anterrieu (1999)), there are strong connections between phase- and phasor-based approaches. To see this, let \mathbf{z} be the vector of products $\{g_i g_j^* f_{|i-j|}\}$ which minimize Λ . We rewrite Equation (3.38) as:

$$\Lambda = \sum_i \sum_{j>i} w_{ij} ||(V_{ij} - z_{ij})(V_{ij}^* - z_{ij}^*)|| \quad (3.41)$$

Define $r_{ij} = e^{j2\pi n_{ij}} z_{ij}$ for an arbitrary integer n_{ij} and \mathbf{r} as the vector containing the r_{ij} .

Note that \mathbf{r} also minimizes Λ since the rotations $\{e^{j2\pi n_{ij}}\}$ do not change the values of the residuals in Λ . Hence any set of rotated phasors $\{\tilde{g}_i\}$ and $\{\tilde{f}_{|i-j|}\}$ whose products

produce the vector \mathbf{r} will also minimize Λ . Note that the set of such valid phase vectors (i.e. the concatenations of possible $\{\angle \tilde{g}_i\}$ and $\{\angle \tilde{f}_{|i-j|}\}$) includes the complete family of Phase-approach solutions τ_σ in Section 3.7.1 with $\hat{\beta}_{K+L}^* = \angle \mathbf{r}$ (where $\angle \mathbf{r}$ is the vector of the phases of the complex vector \mathbf{r}). In other words, the valid phase component of the phasor approach solutions is not unique, and the minimization of Λ admits the same solution ambiguity depicted in Figure 3.6. Hence we see that integer ambiguities present in the phase approach do not disappear in the phasor approach; in fact, the unwrapped candidate solutions of the phase-based approach correspond to minima of the phasor-based objective. In practice Phasor Approach techniques may converge to any one of these minima. Hence the critical issue for reliable image reconstruction is again the nature of the difference between these valid minima and the true solutions. Based on the connections we have drawn with Phase-approach solutions above, the following Proposition is clear:

Proposition 3.7.4. *If a pattern is wrap-invariant, the global minima of the Phasor-Approach objective caused by the inherent 2π -periodicity in the objective's residual differ from the true solution merely by an image shift.*

In practice we see that, as in the Phase approach, if the pattern is not wrap-invariant, the Phasor Approach suffers from false global minima producing severe distortion in the resulting reconstruction. The lower right panel of Figure 3.3 shows the resulting reconstruction produced by the Phasor method using the updates in Equations (3.39) and (3.40). To show the correspondence in solutions between two types of approaches, we provided the result of the Phase approach as the initial point for the Phasor approach as is common in practice (Liu *et al.*, 2010), (Zheng *et al.*, 2014). Indeed this point is a global minimum of Λ which the updates above cannot escape, and as a result we observe virtually-identical distortion to that of the Phase approach.

We repeated the experiment with the wrap-invariant pattern in Figure 3.11, and as expected, the results were pristine in the noiseless case and virtually-identical to those of the Phase approach in the noisy case. For completeness the results are given in Figure 3.13.

As expected, different initializations of the updates in Equations (3.39) and (3.40) simply

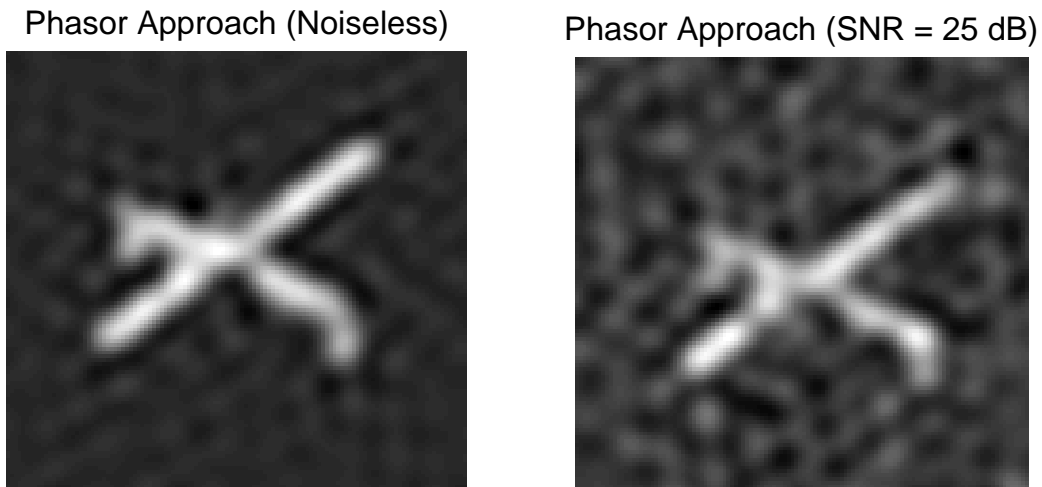


Figure 3.13: *Reconstruction Results for Phasor Approach*

resulted in shifts of an otherwise pristine image in the reconstruction.

3.8 Conclusions

In this chapter, we have examined the ambiguities caused by the 2π -periodicity of interferometric phase in Redundant Spacing Calibration. In particular we have described their fundamental presence in existing RSC methods whether the observables considered are the measured baseline phasors or their phases. In the former, e.g. by Greenaway (1990), they are manifested as phase-wrapping errors, and in the latter, e.g. by Marthi and Chengalur (2014), as false minima. We have demonstrated that in either case, these ambiguities can result in noticeable distortion of the reconstructed image.

Using the Closest-Vector-Problem formulation of the unwrapping problem due to Lannes and Anterrieu (1999), we have developed the notion of a wrap-invariant pattern. For wrap-invariant patterns, the impact of the 2π -periodicity can be completely eliminated (c.f. Section 3.5) using well-known algorithms from lattice theory and the Smith Normal Form, and reduced to a mere image shift (c.f. Section 3.7) when existing, fast approaches are used. Phase-approach solutions (Arnot *et al.*, 1985), (Greenaway, 1990), (Wieringa, 1992), (Lannes and Anterrieu, 1999)) are commonly used to quickly obtain an initial point to aid in the

convergence of the Phasor approach. They are obtained by selecting specific solutions from the general family in Equation (3.25) via enforcement of additional constraints on the solution. Phasor approaches seek those complex gains and object visibilities which minimize a squared-residual with respect to the observed complex visibilities (Wijnholds and Noorishad, 2012), (Liu *et al.*, 2010), (Marthi and Chengalur, 2014) using, for example, gradient-descent methods. We have seen that the same 2π -ambiguity which creates a family of Phase-approach solutions also produces a corresponding family of global minima in the Phasor approach. For either case, our results show that for wrap-invariant patterns, this family represents merely shifted versions of the true image. We have also extended this analysis to show that wrap-invariant patterns admit reliable imaging using standard, and generalized, phase closures. Conversely we show that patterns which are not wrap-invariant can suffer from distortion of the sort depicted in Figure 3.3.

The prognosis for mitigation of the ambiguity issues raised in this chapter is quite positive. Random patterns appear to satisfy the wrap-invariance condition with high probability. Moreover, failure to meet this condition amounts to the existence of a particular kind of cycle in the interferometric graph which can be easily isolated; a chief contribution of this chapter is a simple algorithm for identifying such cycles so that they can be removed by the array designer. Finally, we have shown an example execution of the algorithm to diagnose a member of a popular pattern which is not wrap-invariant. It is clear that with careful array design, both Phase- and Phasor-based RSC techniques can reliably produce quality image reconstructions free from discernible artifacts.

Chapter 4

Robust image reconstruction with redundant arrays and generalized closure phases

4.1 Chapter Overview

Redundant Spacing Calibration (RSC) techniques employ redundancy in the baselines of a telescope array to eliminate the contribution of atmospheric turbulence in the interferometric observables. Whereas conventional techniques for this elimination require the enforcement of prior constraints on the underlying image, RSC algorithms can be, in principle, mathematically well-posed and hence require no such prior knowledge. Traditionally these algorithms have been applied directly to the fringe measurements. However, in scenarios of low photon flux, such as those arising in the observation of dim objects in space, single-exposure fringe measurements are not reliable observables in general. Instead one must rely on time-averaged, atmosphere-invariant quantities such as the bispectrum. In this chapter, we develop a novel algorithm for redundant arrays which provides robust image reconstruction using integrable atmosphere-invariant observables. Our algorithm utilizes standard linear estimation methods, as well as techniques from lattice theory, to reliably estimate the Fourier

phase. Moreover, we provide theoretical and empirical evidence that generalizing the classical bispectrum to higher-order atmosphere-invariant observables, which we call *n-spectra*, can offer significant performance gains. Our selection of an independent and high-SNR set of *n-spectra* leverages the notion of the *minimum cycle basis* from graph theory. We analyze the expected shot-noise-limited performance of our algorithm for both pairwise and Fizeau interferometric architectures, and corroborate this analysis with simulation results showing performance near the Cramér-Rao bound. Lastly, we apply techniques from the field of *sparse recovery* (SR) to perform image reconstruction from the estimated complex visibilities.

4.2 Problem Statement and Prior Work

Recall from previous Chapters that interferometric measurements take the form:

$$\vec{\beta} = \vec{\theta} + \mathbf{A}\vec{\phi} \quad (4.1)$$

where \mathbf{A} is an $\binom{N}{2} \times N$ matrix whose rows compute the piston differences involved in each measurement.

In Chapter 2 we introduced a class of algorithms which decouple the Fourier phases $\vec{\theta}$ from the atmosphere contributions $\vec{\phi}$ by imposing prior constraints. With estimates for sparse samples of the Fourier Transform, we showed that the residual ill-posed inverse problem could be regularized effectively using techniques from sparse recovery. As interferometric systems are tasked with the observation of increasingly-complex scenes, approaches which minimize reliance on specialized prior assumptions are expected to gain appeal. In Chapter 3 we described the family of RSC techniques as an intrinsically well-posed alternative to prior-regularized phase recovery is to use baseline redundancy to explicitly solve for piston variation. We also examined the important issue of integer phase ambiguities in the context of these techniques. Array design considerations based on the notion of *wrap-invariance* (Kurien *et al.*, 2016), as well as algorithms from lattice theory (Lannes and Anterrieu, 1999), can be used to solve these ambiguity issues.

In this chapter, we leverage elements from both Chapters 2 and 3 to develop a novel algorithmic framework for low-flux imaging scenarios (as in, for example, observation of dim objects in space). Here we must again rely on atmosphere-invariant observables like the bispectrum and closure phase. Many approaches have been developed for estimation of the Fourier phase from the closure phase relations. Several early papers on the subject (Rogstad, 1968), (Bartelt *et al.*, 1984) proposed solving these equations recursively. Namely, two initial Fourier phases $\hat{\theta}_A$ and $\hat{\theta}_B$ associated with two connected baselines A and B can be initialized arbitrarily without loss of generality (see Section 3.5.1). The Fourier phase $\hat{\theta}_C$ of the residual baseline C closing A and B can be determined as: $\hat{\theta}_C = \hat{\beta}_{ABC} - \hat{\theta}_A - \hat{\theta}_B$, where $\hat{\beta}_{ABC}$ is the closure phase associated with triangle ABC . These initial phase estimates can then *bootstrap* solution for the remaining phases using the appropriate closure relations involving redundant copies of the baselines associated with these initial phases. While simple to implement, such approaches are clearly sub-optimal in terms of estimation accuracy; the fact that information about a given Fourier phase is contained in multiple closure phases implies there is a benefit in joint inference of the Fourier phases from the complete set of closure phases.

As we will quantify later in this Section, closure phases are corrupted with a variable amount of measurement noise; closure phases involving short-baseline measurements will typically incur less distortion due their high visibility (and hence high-SNR) relative to their long-baseline counterparts. Reliable joint inference from the closure phases then entails minimization of a weighted sum of closure or bispectrum residuals (i.e. weighted least-squares methods). Such methods are of course complicated by wrapping issues analogous to those described in Chapter 3. To address these, Haniff (1991) developed an algorithm which respects the 2π -periodicity of the closure phase. This approach uses a conjugate-gradient routine to find the set of Fourier phases $\{\theta\}$ which minimizes the following non-linear objective:

$$\Psi_1 = \sum_{i=1}^{N_c} \left(\frac{\text{mod } 2\pi \{ \hat{\beta}_{i,cl} - (\theta_{i1} + \theta_{i2} + \theta_{i3}) \}}{w_i} \right)^2 \quad (4.2)$$

where N_c is the number of closures, $\hat{\beta}_{i,cl}$ is the i -th closure phase, the θ_{i1} , θ_{i2} , and θ_{i3} are the Fourier phases associated with this closure, and w_i is a weighting factor proportional to the estimated variance of the closure measurement.

While this minimization has been shown to be successful in certain cases, it is prone to stagnation at local minima due to the non-smoothness of the $\bmod(x)$ function (Negrete-Regagnon, 1996),(Thiébaud and Giovannelli, 2010). An alternative least-squares approach suggested by Gorham *et al.* (1989) which instead works with the closure phasor (i.e. the normalized bispectrum) has hence generally been preferred (Negrete-Regagnon, 1996). This approach replaces the objective above with one taking the form:

$$\Psi_2 = \sum_{i=1}^{N_c} \frac{\left\| e^{j\hat{\beta}_{i,cl}} - e^{j(\theta_{i1} + \theta_{i2} + \theta_{i3})} \right\|^2}{w_i^2} \quad (4.3)$$

Objective Ψ_2 is clearly a continuous function and hence local minima can be reliably found with gradient-descent techniques. However the objective is non-convex in general, and hence gradient-based techniques are not guaranteed to converge to the true global minimum.

To the best of our knowledge, work by Lannes and Anterrieu (1999) was the first to rigorously study the use of closure phases within the RSC framework. In this work, the authors present a mixed, integer linear-least squares formulation of the phase estimation problem in which the integer component pertains to the unknown (integer) phase wraps. They first propose the use of techniques from lattice theory to resolve the wrap ambiguities. They then develop a pseudo-inverse estimator which recovers the Fourier phases from the measured baseline phases via an intermediate computation of the closure phases. The computation of the necessary covariance matrices for optimal estimation within this formulation is omitted. Since it is in fact the closure phases which are the direct observables in practical low-flux scenarios, we present instead an approach working solely with the closure phases and compute the necessary covariances in straightforward manner. Lannes (2003) later proposed an alternate estimator based on computation of the Smith Normal Form of the matrix mapping Fourier phases to closure phases. This estimator has the advantage of completely eliminating the effect of phase wrapping in the closure measurements when

appropriate routines from lattice theory are used to pre-process the phase measurements. In contrast, our estimator reliably uses standard, fast linear estimation techniques, thereby obviating computation of the Smith Normal Form at the possible expense of an extra shift in the recovered image. Such image shifts are anticipated to be of negligible importance in most cases of practical interest.

Though the bispectrum and associated closure phase have been standard interferometric observables for decades, it is not difficult to imagine situations in which use of these observables unnecessarily limits reconstruction performance. In Figure 4.1, we illustrate this idea with an example. We label each baseline of an interferometric array with a visibility, which is an indicator of the strength of its Fourier component relative to the overall brightness of the image. In accordance with the power-law decay of intensity with spatial-frequency modulus in an overwhelmingly-large fraction of natural images (see, e.g. Ruderman (1994)), we consider a visibility distribution which drops sharply with baseline length. In standard bi-spectral imaging, the high-spatial-frequency phase information associated with long baseline \mathbf{b}^* is recovered via forming closures, for example, with a short baseline along with another long baseline (*Traditional*). As we will see later in the chapter, for a certain range of photon fluxes, the SNR of the bispectrum is approximately proportional to the sum of the reciprocals of the squared visibilities of the associated baselines (Kulkarni *et al.*, 1991), i.e.:

$$SNR \approx \hat{n} \left(\sum_{i=1}^3 \frac{1}{\gamma_i^2} \right)^{-1} \quad (4.4)$$

where \hat{n} is the number of photoelectrons (pe) received per interference fringe per exposure frame. As we will show in this chapter, this SNR model extends in the natural way to the SNR of higher-order observables like that shown on the right (*Generalized*). In contrast with the *Traditional* observable, *Generalized* observable utilizes only short (high-visibility) baselines to close the long baseline, resulting in a near-doubling of the resulting SNR. We will henceforth refer to the higher-order generalization of the bispectrum triple product as the *n-spectrum*, and its phase as the *generalized closure phase*.

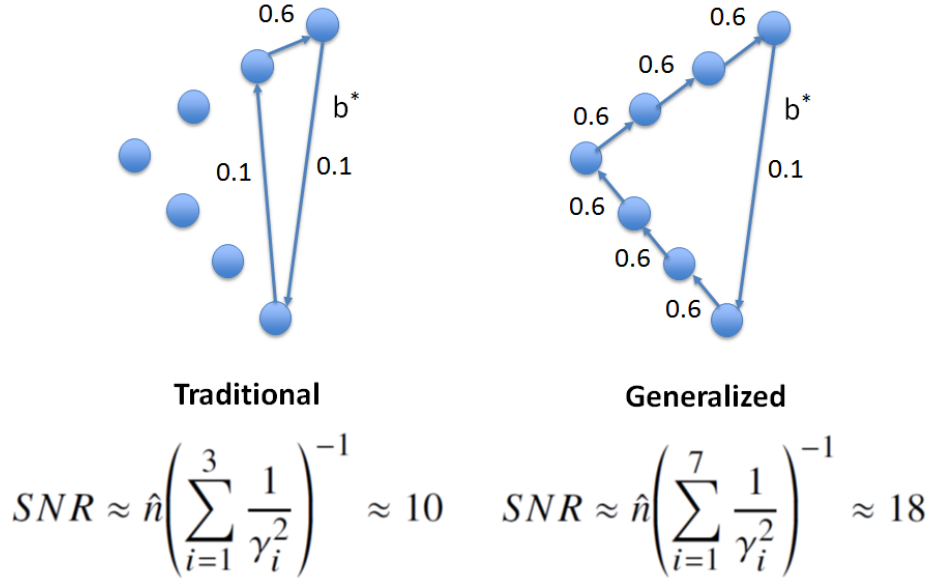


Figure 4.1: Generalizing the phase closure concept . SNRs given assume each aperture contributes $\hat{n} = 2e3$ photoelectrons to each fringe in pairwise combination.

Generalization of the closure phase is not a novel concept in the astronomical community. Recently Martinache (2010) proposed observables known as *Kernel phases*, which are obtained by computing an orthogonal basis for the left-nullspace (e.g. using the Singular Value Decomposition) of the matrix \mathbf{A} in Equation (4.1) . The vectors of this basis are assembled as the rows of a matrix \mathbf{K} . When this matrix is applied to the measurements β , the contributions due to the atmosphere (i.e. $\vec{\phi}$) clearly cancel, leaving only linear combinations of the object phases in θ . These linear combinations can then be used to reconstruct the Fourier phase via direct application of the pseudo-inverse of \mathbf{K} (Martinache, 2014), or via regularized inversion using prior image models (Ireland, 2013). These studies consider masked-aperture or segmented-pupil scenarios in which either closure or Kernel phases can be obtained directly via linear combinations of the raw phase measurements. Since we consider the scenario encountered in long-baseline interferometry, in which accurate raw phase measurements are not reliably obtained via multi-frame averaging, our inference problem must be formulated in terms of integrated phasor observables, i.e. the *n-spectra* defined above.

Our approach is also conceptually similar to the approach of *baseline bootstrapping* introduced in Chapter 3, Section 3.5.3. To see this, let us define the normalized *optical path difference (OPD)* as the difference in the atmosphere phases between a given pair of apertures. In the bootstrapping approach first proposed by Roddier (1988), temporal OPD changes along a long, low-SNR baseline \mathbf{b} are inferred by combining (or *bootstrapping*) those measured on several short high-SNR baselines along a closed path involving \mathbf{b} . In principle, the phases on all baselines can then be de-rotated accordingly and the corresponding phasors integrated coherently, thereby boosting the SNR on all baselines. This tracking and de-rotation can be done real-time with fringe tracking or offline on processed data (Buscher, 2015). As in bootstrapping, our approach seeks to build reliable proxy observables for the phase estimates of the low-visibility baselines by forming closed paths involving high-visibility baselines. However, bootstrapping eliminates differential phase relative to a reference exposure frame and hence implicitly requires a threshold per-frame SNR on the shortest baselines for robustness. In contrast, our approach has no such threshold as it works directly with atmosphere-invariant observables.

In this chapter, we develop a comprehensive algorithmic framework for image reconstruction based on RSC techniques and sparse recovery, and analyze its performance for the two popular beam-combining architectures used in optical interferometry (*pairwise*, and *Fizeau*). The chapter is organized as follows. In Section II, we derive an SNR model for the n -spectrum, as well as for the covariance amongst distinct n -spectra. The analysis begins with the pairwise case for simplicity, and mathematical arguments for an extension to the Fizeau case are presented at the end of the section and in the Appendix. In Section III, we describe a systematic, integer-least-squares approach for reconstructing imagery from object visibilities and generalized closure phases (GC's). We draw on the notion of a *minimum cycle basis* to select a minimum set of linearly-independent GC's (i.e. a *basis*) of minimum variance which spans the subspace of all closures. Leveraging techniques first suggested by Lannes and Anterrieu (1999) and the notion of *wrap-invariant* measurement mappings, we apply algorithms from lattice theory to reliably unwrap the GC's. We then quantify the gain

in theoretical performance afforded by generalizing the classical third-order atmosphere-invariant observables to higher-order cycles, and corroborate this theoretical performance with simulation results. In this simulation, we apply our algorithm to reconstruct a dim, structured object from generalized closures. For the final image reconstruction, we employ a sparse recovery algorithm which enforces smoothness in the image domain by minimizing a *total variation* metric.

4.3 Preliminaries

In this chapter we analyze the performance of generalized-closure-based imaging for both the pairwise and Fizeau interferometric architectures (see Section 1.2) in terms of mean-squared error (MSE) of the phase estimates. We begin with the pairwise architecture since it is significantly easier to analyze. We then present a series of approximations which accurately describe the algorithm's performance in the Fizeau case while keeping the tediousness of the mathematics at a manageable level. In both cases, the analysis is based on the computation of the moments of the Poisson distribution; in the pairwise case, the first two moments are required, whereas in the Fizeau case, the first o moments are needed, where o is the order of the generalized closure.

4.3.1 Fringe Noise Model for Pairwise Beam Combiner

Given N_{ap} apertures, recall that the Fourier phase and amplitude of the object is encoded in a series of $\binom{N_{ap}}{2}$ interference fringes observed on a focal plane. In the pairwise case, each fringe is measured on a separate set of detectors so that the Poisson shot noise incurred is independent for each fringe. Suppose each aperture receives n photons and that this light is split evenly (e.g. by beam-splitter) before being combined with the light from the other apertures. Hence each aperture contributes $\hat{n} = \frac{n}{N_{ap}-1}$ photons to each fringe. The expectation of number of photons q at a given detector k for a given fringe is given by:

$$\langle q(k) \rangle = \frac{2\hat{n}}{N_p} (1 + \gamma \cos(\omega k + \theta + \phi)) \quad (4.5)$$

where N_p is the number of pixels over which these photons are spread, γ is the visibility of the fringe (which takes its value between 0 and 1), ω is the fringe frequency, θ is the Fourier phase, and ϕ is the atmosphere phase.

To retrieve the object phase and amplitude from its fringe encoding, we simply take the Fourier transform and evaluate it at the fringe frequency:

$$z = \sum_{k=0}^{N_p-1} q(k) e^{-i\omega k} \quad (4.6)$$

The expectation of this quantity is given by:

$$\langle z \rangle = \hat{n} \gamma e^{-i\beta_f} \quad (4.7)$$

where β_f is the phase of the fringe.

To obtain the bispectrum SNR, it will be useful to compute the power in the fringe, i.e. expectation of the quantity zz^* :

$$\langle zz^* \rangle = \left\langle \sum_{k=0}^{N_p-1} \sum_{k'=0}^{N_p-1} q(k) q(k') e^{-i\omega k} e^{i\omega k'} \right\rangle \quad (4.8)$$

This is equivalent to:

$$\langle zz^* \rangle = \sum_{k=0}^{N_p-1} \sum_{k'=0}^{N_p-1} \langle q(k) q(k') \rangle e^{-i\omega k} e^{i\omega k'} \quad (4.9)$$

Using the first two moments of the Poisson distribution, we can write:

$$\langle q(k) q(k') \rangle = \langle q(k) \rangle \langle q(k') \rangle + \delta_{kk'} \langle q(k) \rangle \quad (4.10)$$

Substituting into Equation (4.9), we obtain:

$$\langle zz^* \rangle = \sum_{k=0}^{N_p-1} \langle q(k) \rangle e^{-i\omega k} \sum_{k'=0}^{N_p-1} \langle q(k') \rangle e^{i\omega k'} + \sum_{k=0}^{N_p-1} \langle q(k) \rangle \quad (4.11)$$

Simplifying we obtain:

$$\langle zz^* \rangle = \gamma^2 \hat{n}^2 + 2\hat{n} \quad (4.12)$$

4.3.2 N-Spectrum Covariance and SNR Models for the Pairwise Architecture

Let us define a *generalized N-spectrum* G as the product of complex baseline measurements along an o -edge cycle. Given variances $\sigma_{Re}^2(G)$ and $\sigma_{Im}^2(G)$ of real and imaginary components, respectively, we define the pseudo-variance of the product as:

$$V_{pairwise}(G) := \sigma_{Re}^2 + \sigma_{Im}^2 = \langle GG^* \rangle - \langle G \rangle \langle G^* \rangle \quad (4.13)$$

or:

$$V_{pairwise}(G) = \left[\prod_{i=1}^o \langle z_i z_i^* \rangle \right] - \prod_{i=1}^o \langle z_i \rangle \langle z_i^* \rangle \quad (4.14)$$

$$V_{pairwise}(G) = \left[\prod_{i=1}^o (\hat{n}^2 \gamma_i^2 + 2\hat{n}) \right] - \hat{n}^{2o} \prod_{i=1}^o \gamma_i^2 \quad (4.15)$$

In the specific case of the classic triple product (i.e. the *bispectrum*), this expands to (Kulkarni *et al.*, 1991):

$$V_{pairwise}(G) = 2\hat{n}^5(\gamma_1^2\gamma_2^2 + \gamma_2^2\gamma_3^2 + \gamma_3^2\gamma_1^2) + 4\hat{n}^4(\gamma_1^2 + \gamma_2^2 + \gamma_3^2) + 8\hat{n}^3 \quad (4.16)$$

The SNR $\mathcal{S}_{pairwise} := \frac{\sqrt{2}|G|}{\sqrt{V(G)}}$ is then given by:

$$\mathcal{S}_{pairwise} = \frac{\gamma_1\gamma_2\gamma_3\hat{n}^{\frac{3}{2}}}{\sqrt{\hat{n}^2(\gamma_1^2\gamma_2^2 + \gamma_2^2\gamma_3^2 + \gamma_3^2\gamma_1^2) + 2\hat{n}(\gamma_1^2 + \gamma_2^2 + \gamma_3^2) + 4}} \quad (4.17)$$

To gain insight into the limiting behavior of the SNR, we make two approximations: A low-SNR approximation can be found by considering only the constant term in the denominator, which will dominate when \hat{n} and/or the γ_i are small.

$$\mathcal{S}_{pairwise,low} \approx \frac{\gamma_1\gamma_2\gamma_3\hat{n}^{\frac{3}{2}}}{2} \quad (4.18)$$

A high-SNR approximation can be found by considering the highest-order term in the denominator, which will dominate when \hat{n} is large and/or the γ_i are large.

$$\mathcal{S}_{pairwise,high} \approx \frac{\gamma_1\gamma_2\gamma_3\hat{n}^{\frac{3}{2}}}{\sqrt{\hat{n}^2(\gamma_1^2\gamma_2^2 + \gamma_2^2\gamma_3^2 + \gamma_3^2\gamma_1^2)}} \quad (4.19)$$

By averaging the fringe phasors over a sufficient number of frames N_f , we can build the SNR to a level at which the bispectrum phase (i.e. the *closure*) phase can be measured. For phasor SNRs sufficiently greater than 1, we can apply the following well-known approximation for the variance on the corresponding phase θ_{cl} (see e.g. Walkup and Goodman (1973)):

$$\sigma_{\theta_{cl}}^2 \approx \frac{1}{N_f S^2} \quad (4.20)$$

We then obtain the following approximations for the closure phase variances. At low SNR, we have:

$$\sigma_{\theta_{cl},low}^2 \approx \frac{4}{N_f \gamma_1^2 \gamma_2^2 \gamma_3^2 \hat{n}^3} \quad (4.21)$$

Taking the logarithm we find that the variance decouples:

$$\log \sigma_{\theta_{cl},low}^2 \approx -2 \log \gamma_1 - 2 \log \gamma_2 - 2 \log \gamma_3 - 3 \log \hat{n} + C \quad (4.22)$$

where C is a constant independent of the flux level \hat{n} and the visibilities γ_i .

At high SNR, we have:

$$\sigma_{\theta_{cl},high}^2 \approx \frac{1}{\hat{n} N_f} \left(\frac{1}{\gamma_1^2} + \frac{1}{\gamma_2^2} + \frac{1}{\gamma_3^2} \right) \quad (4.23)$$

Generalizing the results above to n -spectra, we obtain for low SNR,

$$\sigma_{\theta_{cl},low}^2 \approx \frac{2^o}{N_f \hat{n}^o \prod_{i=1}^o \gamma_i^2} \quad (4.24)$$

and for high SNR,

$$\sigma_{\theta_{cl},high}^2 \approx \frac{1}{\hat{n} N_f} \sum_{i=1}^o \frac{1}{\gamma_i^2} \quad (4.25)$$

where o is the order of the generalized closure. Note that if we take the logarithm of Equation (4.26), we obtain:

$$\log \sigma_{\theta_{cl},low}^2 \approx o \log \frac{2}{\hat{n}} - 2 \sum_{i=1}^o \log \gamma_i - \log N_f \quad (4.26)$$

4.3.3 Covariance Matrix of the Generalized Closure Phases for the Pairwise Architecture

In this section we generalize results of Kulkarni *et al.* (1991) to obtain expressions for the n-spectra covariance matrices. We will first compute the covariance between two n-spectrum phasors $\mathbf{G}_1 = \text{Re}[G_1] + j\text{Im}[G_1]$, with $E[\mathbf{G}_1] = \mathbf{g}_1$, and $\mathbf{G}_2 = \text{Re}[G_2] + j\text{Im}[G_2]$, and $E[\mathbf{G}_2] = \mathbf{g}_2$. Let I and J denote the set of baselines in n-spectra G_1 and G_2 , respectively. Then we have:

$$\sigma(G_1, G_2^*) = \left\langle \prod_{i \in I} z_i \prod_{j \in J} z_j^* \right\rangle - \prod_{i \in I} \langle z_i \rangle \prod_{j \in J} \langle z_j^* \rangle \quad (4.27)$$

We can then factor out the factors common to G_1 and G_2 from those which are distinct to obtain:

$$\sigma(G_1, G_2^*) = \left[\prod_{k \in I \cap J} \langle z_k z_k^* \rangle - \prod_{k \in I \cap J} \langle z_k \rangle \langle z_k^* \rangle \right] \prod_{l \in I \setminus J} \langle z_l \rangle \prod_{m \in J \setminus I} \langle z_m^* \rangle \quad (4.28)$$

Making the appropriate substitutions from Equation (4.12), we obtain:

$$\begin{aligned} \sigma(G_1, G_2^*) = e^{2\pi j(\angle G_1 - \angle G_2)} & \left[\prod_{k \in I \cap J} (\hat{n}^2 \gamma_k^2 + 2\hat{n} \delta_{z_1, k, z_2, k}) - \prod_{k \in I \cap J} \hat{n}^2 \gamma_k^2 \right] \\ & \times \prod_{l \in I \setminus J} \hat{n} \gamma_l \prod_{m \in J \setminus I} \hat{n} \gamma_m \quad (4.29) \end{aligned}$$

where $\delta_{z_1, k, z_2, k} = 1$ if the common measurements are the same, and zero if they are conjugates of each other. Similarly,

$$\begin{aligned} \sigma(G_1, G_2) = e^{2\pi j(\angle G_1 + \angle G_2)} & \left[\prod_{k \in I \cap J} (\hat{n}^2 \gamma_k^2 + 2\hat{n} \bar{\delta}_{z_1, k, z_2, k}) - \prod_{k \in I \cap J} \hat{n}^2 \gamma_k^2 \right] \\ & \times \prod_{l \in I \setminus J} \hat{n} \gamma_l \prod_{m \in J \setminus I} \hat{n} \gamma_m \quad (4.30) \end{aligned}$$

where $\bar{\delta}_{z_1, k, z_2, k} = 1$ if the common measurements are conjugates of each other, and zero if

they are the same.

We can now compute real and imaginary n-spectra covariance as:

$$\sigma(\text{Re}[G_1], \text{Re}[G_2]) = \frac{1}{2} \text{Re}[\text{Cov}(G_1, G_2^*) + \text{Cov}(G_1, G_2)] \quad (4.31)$$

$$\sigma(\text{Im}[G_1], \text{Re}[G_2]) = \frac{1}{2} \text{Im}[\text{Cov}(G_1, G_2^*) + \text{Cov}(G_1, G_2)] \quad (4.32)$$

$$\sigma(\text{Re}[G_1], \text{Im}[G_2]) = \frac{1}{2} \text{Im}[-\text{Cov}(G_1, G_2^*) + \text{Cov}(G_1, G_2)] \quad (4.33)$$

$$\sigma(\text{Re}[G_1], \text{Im}[G_2]) = \frac{1}{2} \text{Re}[-\text{Cov}(G_1, G_2^*) - \text{Cov}(G_1, G_2)] \quad (4.34)$$

Given the n-spectra covariance expression above, we can now approximate the covariance between two generalized closure phases using Taylor series approximation of the moments, which is a technique known as the *delta method* (see, e.g. Casella and Berger (2002)). First we define the angle function of a phasor as $\theta(x, y) = \arctan \frac{y}{x}$. The corresponding generalized closure phases are given by: $\Theta_1 = \theta(\text{Im}[G_1], \text{Re}[G_1])$ and $\Theta_2 = \theta(\text{Im}[G_2], \text{Re}[G_2])$. Applying the delta method yields the following approximation:

$$\sigma(\Theta_1, \Theta_2) \approx \left(\frac{\partial \theta}{\partial x} \Big|_{\mathbf{g}_1} \right) \left(\frac{\partial \theta}{\partial y} \Big|_{\mathbf{g}_2} \right) \sigma(\text{Re}[G_1], \text{Im}[G_2]) \quad (4.35)$$

$$+ \left(\frac{\partial \theta}{\partial y} \Big|_{\mathbf{g}_1} \right) \left(\frac{\partial \theta}{\partial x} \Big|_{\mathbf{g}_2} \right) \sigma(\text{Im}[G_1], \text{Re}[G_2]) \quad (4.36)$$

$$+ \left(\frac{\partial \theta}{\partial y} \Big|_{\mathbf{g}_1} \right) \left(\frac{\partial \theta}{\partial y} \Big|_{\mathbf{g}_2} \right) \sigma(\text{Im}[G_1], \text{Im}[G_2]) \quad (4.37)$$

$$+ \left(\frac{\partial \theta}{\partial x} \Big|_{\mathbf{g}_1} \right) \left(\frac{\partial \theta}{\partial x} \Big|_{\mathbf{g}_2} \right) \sigma(\text{Re}[G_1], \text{Re}[G_2]) \quad (4.38)$$

Evaluation of the derivatives in the expression above yields:

$$\begin{aligned}
\sigma(\Theta_1, \Theta_2) \approx & \left(\frac{-Im[g_1]}{Re[g_1]^2 + Im[g_1]^2} \right) \left(\frac{Re[g_2]}{Re[g_2]^2 + Im[g_2]^2} \right) \sigma(Re[G_1], Im[G_2]) \\
& + \left(\frac{Re[g_1]}{Re[g_1]^2 + Im[g_1]^2} \right) \left(\frac{-Im[g_2]}{Re[g_2]^2 + Im[g_2]^2} \right) \sigma(Im[G_1], Re[G_2]) \\
& + \left(\frac{Re[g_1]}{Re[g_1]^2 + Im[g_1]^2} \right) \left(\frac{Re[g_2]}{Re[g_2]^2 + Im[g_2]^2} \right) \sigma(Im[G_1], Im[G_2]) \\
& + \left(\frac{-Im[g_1]}{Re[g_1]^2 + Im[g_1]^2} \right) \left(\frac{-Im[g_2]}{Re[g_2]^2 + Im[g_2]^2} \right) \sigma(Re[G_1], Re[G_2]) \quad (4.39)
\end{aligned}$$

4.3.4 Covariance Matrix Approximations for the Fizeau Architecture

In the Fizeau case, Equation (4.12) becomes:

$$\langle zz^* \rangle = \gamma^2 n^2 + N_{ap} n \quad (4.40)$$

As shown in Appendix C.1, we can approximate the variance of a n-spectrum of order o as:

$$V_{Fizeau}(G) \approx \left[\prod_{i=1}^o (n^2 \gamma_i^2 + N_{ap} n) \right] - n^{2o} \prod_{i=1}^o \gamma_i^2 \quad (4.41)$$

Hence the SNR can be approximated as:

$$\mathcal{S}_{Fizeau} \approx \frac{\sqrt{2} n^{\frac{o}{2}} \prod_{i=1}^o \gamma_i}{\sqrt{\left[\prod_{i=1}^o (n^2 \gamma_i^2 + N_{ap} n) \right] - n^{2o} \prod_{i=1}^o \gamma_i^2}} \quad (4.42)$$

Applying the same approximations in Appendix C.1 to the covariance yields the following analogues to Equations (4.29)-(4.30)

$$\begin{aligned}
\sigma(G_1, G_2^*) = & e^{2\pi j(\angle G_1 - \angle G_2)} \\
& \times \left[\prod_{k \in I \cap J} (n^2 \gamma_k^2 + N_{ap} n \delta_{z_{1,k}, z_{2,k}}) - \prod_{k \in I \cap J} n^2 \gamma_k^2 \right] \prod_{l \in I \setminus J} n \gamma_l \prod_{m \in J \setminus I} n \gamma_m \quad (4.43)
\end{aligned}$$

and,

$$\sigma(G_1, G_2) = e^{2\pi j(\angle G_1 + \angle G_2)} \times \left[\prod_{k \in I \cap J} (n^2 \gamma_k^2 + N_{ap} n \bar{\delta}_{z_{1,k}, z_{2,k}}) - \prod_{k \in I \cap J} n^2 \gamma_k^2 \right] \times \prod_{l \in I \setminus J} n \gamma_l \prod_{m \in J \setminus I} n \gamma_m \quad (4.44)$$

4.4 Fourier Phase Recovery using the N-Spectra

In this section, we briefly review the integer least-squares approach for RSC phase recovery (see e.g. Lannes and Anterrieu (1999), Kurien *et al.* (2016) for more comprehensive treatments), and then describe an algorithm for selection of a set of generalized closure relations of minimum total variance.

4.4.1 RSC Phase Recovery with Wrap-Invariant Closure Mappings

We begin by recalling the fundamental RSC equation from Chapter 3:

$$\mathbf{M} \begin{bmatrix} \vec{\theta} \\ \vec{\phi} \end{bmatrix} = \boldsymbol{\beta} + 2\pi \mathbf{e} \quad (4.45)$$

In this Section we discuss generalized closure relations which eliminate the $\vec{\phi}$ contribution in the model in Equation (4.45). We first recall from Chapter 3 (see Definition 3.5.1) that following the analysis of Lannes and Anterrieu (1999), the range of the matrix \mathbf{M} can be divided into two subspaces: (1) the subspace K spanned by the d columns associated with the Fourier phases, or equivalently the range of the sub-matrix \mathbf{M}_θ formed by these columns, and (2) the piston, or *aberration*, phase space L is the span of the N_{ap} columns associated with the piston phases, or equivalently the range of the sub-matrix \mathbf{M}_ϕ formed by these columns. Note that the d columns spanning K are linearly-independent by virtue of having non-zero entries in mutually disjoint sets of baseline indices. Hence $\dim(K) = d$. The subspace L has dimension $N_{ap} - 1$ (Lannes and Anterrieu, 1999); the constant vector forms the one-dimensional nullspace of \mathbf{M}_ϕ .

Additionally, let us establish the following Definitions:

Definition 4.4.1. (Interferometric graph): The interferometric graph of an array is the directed graph whose vertices are the N_{ap} apertures in the array and edges are the $\binom{N_{ap}}{2}$ baselines connecting all vertex pairs.

Definition 4.4.2. (Cycle-space of a directed graph) (Liebchen and Rizzi, 2005): Given a directed graph G with a set of vertices V and edges E , the cycle space of the graph is the vector space of $\mathbb{Q}^{|E|}$ spanned by the incidence vectors of cycles with a (clockwise, or counter-clockwise) *orientation*. It can be shown that the dimension of the cycle space is $M + N - 1$, where $M = |E|$ and $N = |V|$.

Suppose we have a basis for the cycle-space of the interferometric graph. Let us stack the elements of this basis as row vectors in a $\binom{N_{ap}-1}{2} \times \binom{N_{ap}}{2}$ matrix \mathbf{C}_{mc} , which maps baseline phase measurements to generalized closure phases. Recall that closure relations eliminate piston differences in the measurements so that \mathbf{C}_{mc} annihilates the subspace L , i.e. the space spanned by the columns of \mathbf{M} corresponding to $\vec{\phi}$.

In Figure 4.2, we show a simple four-aperture interferometric graph and one possible cycle basis for the graph. For this particular example, we can write:

$$\mathbf{C}_{mc} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} \quad (4.46)$$

where the column (i.e. baseline) indexing in the matrix follows the labeling in the Figure. Note that the final triangle cycle (i.e. the one containing baselines 2, 4, and 6, and represented by the row vector $\mathbf{w}^* := (0, 1, 0, 1, 0, -1)$), can be represented as a linear combination of the cycles in the basis: $\mathbf{w}^* = \mathbf{c}_1 + \mathbf{c}_2 - \mathbf{c}_3$, where $\{c_i\}$ are the rows of \mathbf{C}_{mc}

While for purposes of simplicity we have chosen a cycle basis consisting exclusively of three-baseline cycles, in general the elements of a cycle basis can contain any number of edges.

Lemma 4.4.1. *The nullspace of \mathbf{C}_{mc} is L .*

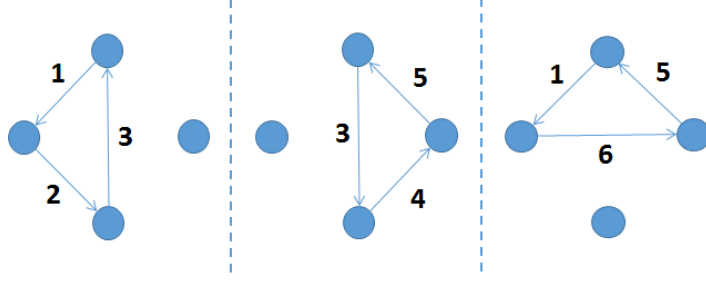


Figure 4.2: One possible cycle-basis for a simple interferometric graph. Note the fourth triangle (i.e. baseline set $\{2, 4, 6\}$) can be expressed as a linear combination of the cycles shown.

Proof: Clearly the N_{ap} columns of L are in the nullspace of \mathbf{C}_{mc} from the arguments above. To see that these columns span the nullspace, note that from the well-known Rank-nullity theorem from linear algebra (see, e.g. Bretscher (2001)), we have:

$$\dim(\ker(\mathbf{C}_{mc})) = \binom{N_{ap}}{2} - \binom{N_{ap}-1}{2} = N_{ap} - 1 \quad (4.47)$$

Hence $\dim(\ker(\mathbf{C}_{mc})) = \dim(L)$. \square

We can now form closure relations as:

$$\mathbf{C}_{mc}\mathbf{M} \begin{bmatrix} \vec{\theta} \\ \vec{\phi} \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{oc} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \vec{\theta} \\ \vec{\phi} \end{bmatrix} \quad (4.48)$$

where $\mathbf{C}_{oc} := \mathbf{C}_{mc}\mathbf{M}_{\theta}$, and is the mapping between object phases and closure phases.

Proposition 4.4.2. Let $\Delta\mathbf{r}_x$ and $\Delta\mathbf{r}_y$ denote the vectors containing the x - and y -coordinates of the baselines in an array, respectively. If the array is a valid RSC array, these columns form a basis for the two-dimensional nullspace of \mathbf{C}_{oc} .

Corollary 4.4.3. For a valid RSC array, the mapping \mathbf{C}_{oc} is injective up to an image shift.

Proofs of Proposition 4.4.2 and Corollary 4.4.3 are given in Appendix C.2. \square

Let us define \mathbf{y}_{cl} as the observation vector of wrapped generalized closure phases, and $2\pi\mathbf{e}_{cl}$ the wrapping vector (with integer \mathbf{e}_{cl}). Our phase measurement model is then given by:

$$\mathbf{C}_{oc}\vec{\theta} = \mathbf{y}_e + \mathbf{n} \quad (4.49)$$

where $\mathbf{y}_e = \mathbf{y}_{cl} + 2\pi\mathbf{e}_{cl}$, and \mathbf{n} is the measurement noise, which we shall assume in this chapter is predominantly due to shot noise on the focal plane.

Given that the generalized-closure mapping is injective up to an image shift, we can now formulate the well-posed recovery of the Fourier phases as the following generalized, integer least-squares problem, which finds a vector in the range of the closure matrix at a minimum Mahalanobis distance from the actual, wrapped measurement vector:

$$\hat{\theta}_{RSC} = \underset{\mathbf{e}_{cl}, \vec{\theta}}{\operatorname{argmin}} \left(\mathbf{y}_e - \mathbf{C}_{oc}\vec{\theta} \right)^T \hat{\Sigma}^{-1} \left(\mathbf{y}_e - \mathbf{C}_{oc}\vec{\theta} \right) \quad (4.50)$$

where $\hat{\Sigma}$ is an estimator of the covariance matrix of the phase measurements. This estimator can be obtained by substituting estimates for the object visibilities $\{\gamma_i\}$ into the covariance expressions in the previous section. It is acknowledged that this least-squares solution is neither the optimal nor maximum-likelihood solution to the phase inference problem. For one, a Gaussian distribution on the closure phases is implicitly assumed to hold, and can only approximately hold for high-SNR closure phases. It is therefore assumed that n-spectra will be integrated for a sufficiently large number of frames for this approximation to be reliable for most, if not all, of the closure phases.

Assuming the covariance matrix estimator is not degenerate, it will admit a Cholesky-decomposition $\Sigma = \mathbf{B}\mathbf{B}^T$. Equation (4.50) is then equivalent to searching for vectors \mathbf{e} which minimize the projection of a whitened measurement $\mathbf{B}^{-1}\mathbf{y}_e$ onto the subspace defined by $\ker((\mathbf{B}^{-1}\mathbf{C}_{oc})^T)$. Specifically we seek to minimize:

$$f(\mathbf{e}) = \|\mathbf{P}_\Sigma \mathbf{B}^{-1}(\mathbf{y}_{cl} + 2\pi\mathbf{e}_{cl})\|^2 \quad (4.51)$$

where \mathbf{P}_Σ is a matrix representing the orthogonal projection from \mathbb{R}^n onto $\ker((\mathbf{B}^{-1}\mathbf{C}_{oc})^T)$.

Letting $\mathbf{e}'_{cl} = -\mathbf{e}_{cl}$, we can rewrite the above objective function as:

$$f(\mathbf{e}') = \|\mathbf{P}_\Sigma \mathbf{B}^{-1}(\mathbf{y}_{cl} - 2\pi\mathbf{e}'_{cl})\|^2 = \|\mathbf{P}_\Sigma \mathbf{B}^{-1}\mathbf{y}_{cl} - 2\pi\mathbf{P}_\Sigma \mathbf{B}^{-1}\mathbf{e}'_{cl}\|^2 \quad (4.52)$$

This optimization problem is equivalent to the so-called *closest vector problem* (CVP) in the theory of lattices. The connection of interferometric phase-unwrapping with the CVP problem was, to the best of our knowledge, first explored in Lannes and Anterrieu (1999). In this paper, the authors formulate an analogous minimization problem for the raw baseline measurements.

We will define a lattice $\mathbf{L}(\mathbb{Z}^n)$ as the set of points generated by integer combinations of the column vectors of a matrix \mathbf{L} . Letting $\tilde{\mathbf{P}} = \mathbf{P}_\Sigma \mathbf{B}^{-1}$, our optimization problem then amounts to the following: Find the lattice point in $\tilde{\mathbf{P}}(\mathbb{Z}^n)$ which is closest to $\tilde{\mathbf{P}}\mathbf{y}_{cl}/(2\pi)$. A compact representation of the lattice Γ is given by:

$$\Gamma = \left\{ \sum_{i=1}^{m \leq n-(d+N-3)} a_i \mathbf{v}_i \mid \forall a_i \in \mathbb{Z} \right\} \quad (4.53)$$

where $\{\mathbf{v}_i\}$ are linearly-independent and together form a *basis* of the lattice. Given the lattice basis, several algorithms exist for finding the closest lattice point to a specified vector. A popular class of algorithms, known as the Sphere-Decoding algorithms, are efficient searches for the closest lattice point within a hypersphere of a certain radius centered on the input vector (see e.g. Agrell *et al.* (2002)). For the simulations in this chapter, we instead use the lower-complexity Babai Nearest Plane (Babai-NP) algorithm (Babai, 1986). For lattice bases which are nearly orthogonal (such as those we use for our simulations), this algorithm offers reliable, albeit not guaranteed, performance in practice.

Suppose we have found a basis for the lattice $\tilde{\mathbf{P}}(\mathbb{Z}^n)$, and we have solved the Closest Vector Problem for a given closure vector \mathbf{y}_{cl} . Let \mathbf{b}^* be the output of the Babai Nearest Plane Algorithm - i.e. it is the lattice point which is the closest to \mathbf{y}_{cl} .¹ We now seek to solve for the wrap vector corresponding to this lattice point, i.e. we seek a solution to:

$$\mathbf{b}^* = \tilde{\mathbf{P}}\mathbf{e} \quad (4.54)$$

¹It is well-known that the performance of the Babai algorithm and other CVP algorithms is improved when the lattice basis is reduced as mentioned in Section 3.3.1. We used the LLL algorithm implementation due to Zhou (2014) in our simulations.

Note that $\tilde{\mathbf{P}}$ is a projection matrix and thus not full-rank, and therefore there will be infinitely-many solutions to this equation. Indeed, by analogy to the integer wrap vectors in Chapter 3 (which were correct only up to an integer vector in the range of \mathbf{M}), our solution \mathbf{y}_{cl}^* is correct only to within an integer vector \mathbf{e}_{resid} in the range of \mathbf{C}_{oc} . Hence there is a fundamental ambiguity that needs to be addressed. Once the measurement vector is unwrapped, there are various ways to solve the RSC system which respect the presence of a residual ambiguity \mathbf{e}_{resid} described above. Following standard least-squares principles, we first compute via projection the vector $\mathbf{y}_{cl,im(C)}^*$ in $im(\mathbf{C}_{oc})$ closest (in Mahalanobis distance) to the unwrapped measurement vector, i.e.

$$\mathbf{y}_{cl,im(C)}^* = \mathbf{B}^{-1}(\mathbf{I} - \mathbf{P}_{\Sigma})\mathbf{B}\mathbf{y}_{cl}^* \quad (4.55)$$

Per the first method (Lannes, 2003), we can compute the Smith Normal Form (SNF) (see Theorems 3.5.2 and 3.5.3) $\{\mathbf{U}, \mathbf{D}, \mathbf{V}\}$ of our matrix \mathbf{C}_{oc} , and set $\mathbf{U}_C = \mathbf{U}^{-1}$, $\mathbf{D}_C = \mathbf{D}$, and $\mathbf{V}_C = \mathbf{V}^{-1}$. We then can write:

$$\mathbf{C}_{oc} = \mathbf{U}_C \mathbf{D}_C \mathbf{V}_C \quad (4.56)$$

A valid RSC solution can then be obtained by evaluating:

$$\hat{\theta}_{RSC} = \mathbf{V}_C^{-1} \mathbf{D}_C^+ \mathbf{U}_C^{-1} \mathbf{y}_{cl,im(C)}^* \quad (4.57)$$

where \mathbf{D}_C^+ denotes the pseudo-inverse of \mathbf{D}_C .

The effect of the residual ambiguity vector \mathbf{e}_{resid} on this solution requires careful consideration. Note that the resulting wrap-induced error in this solution is given by:

$$2\pi \mathbf{e}_{RSC} = 2\pi \mathbf{V}_C^{-1} \mathbf{D}_C^+ \mathbf{U}_C^{-1} \mathbf{e}_{resid} \quad (4.58)$$

Since the matrices \mathbf{U}_C and \mathbf{V}_C are unimodular, they are invertible over the integers. If all the elementary divisors in \mathbf{D}_C are equal to 1, then \mathbf{D}_C^+ will also be integral, and hence \mathbf{e}_{RSC} will be integral. This in turn guarantees that the final error in the Fourier phase will be $\mathbf{0}$

mod 2π , i.e. that the RSC solution $\hat{\theta}_{RSC}$ is wrap-invariant.

Definition 4.4.3. A wrap-invariant closure mapping \mathbf{C}_{oc} is one whose elementary divisors of the matrix \mathbf{C}_{oc} are all 1.

Proposition 4.4.4. (Sufficient condition for wrap-invariant closure imaging): *Given a wrap-invariant closure mapping \mathbf{C}_{oc} , the RSC phase solutions derived from the associated (generalized) closure phases will be immune to the effects of phase-wrapping.*

Recall that in Chapter 3, we observed that this condition can be violated by RSC patterns which are not wrap-invariant. Conversely, we have observed that with wrap-invariant patterns, the closure mappings associated with the cycle bases that we will consider in this paper (i.e. fundamental cycle bases and minimum cycle bases) are also typically wrap-invariant in practice.

As it turns out, given a wrap-invariant closure mapping, the inverse problem in Equation (4.49) can be solved reliably even without the Smith Normal Form; it is sufficient to find any unimodular r -by- r sub-matrix $\tilde{\mathbf{C}}$ of \mathbf{C}_{oc} and solve the associated smaller system in Equation (4.59) to obtain a valid solution. Recall that for valid RSC arrays \mathbf{C}_{oc} will have $r + 2$ columns and hence in this approach, two (non-collinear) object phases are implicitly set to zero. This implicit selection then fixes the fundamentally-ambiguous translation of the scene discussed at the beginning of the section.

$$\tilde{\theta} = \tilde{\mathbf{C}}^{-1} \mathbf{y}_{cl,im(C)}^* \quad (4.59)$$

Since $\tilde{\mathbf{C}}$ is unimodular, $\tilde{\mathbf{C}}^{-1}$ will have solely integral entries so that the resulting wrap error $\tilde{\mathbf{C}}^{-1} \mathbf{e}_{resid} = \mathbf{0} \mod 2\pi$. Assuming the measured phase has been correctly unwrapped, the solution in Equation (4.59) amounts to a so-called *basic solution* of our generalized least-squares problem. We then have the following expression for the error covariance matrix for the estimator $\hat{\theta}_{RSC}$ (Kay, 1993).

$$\Sigma_{basic} = (\tilde{\mathbf{C}}^T \Sigma^{-1} \tilde{\mathbf{C}})^{-1} \quad (4.60)$$

The basic solutions belong to the countably-infinite set of solutions to Equation (4.49). It turns out, somewhat surprisingly, that *any* particular solution within this set can be reliably limited to a mere image shift for patterns that are wrap-invariant. The complete set is given by:

$$\boldsymbol{\theta}_\sigma = \mathbf{C}_{oc}^+ \mathbf{y}_{cl,im(C)}^* + \boldsymbol{\theta}_0 \quad (4.61)$$

where $\boldsymbol{\theta}_0$ is any vector in the nullspace of \mathbf{C}_{oc} , and \mathbf{C}_{oc}^+ denotes the pseudo-inverse of \mathbf{C}_{oc} . The pseudo-inverse can be computed using the Singular-Value-Decomposition (SVD) of \mathbf{C}_{oc} , which is given by:

$$\mathbf{C}_{oc} = \mathbf{U}_\sigma \boldsymbol{\Sigma}_\sigma \mathbf{V}_\sigma^T \quad (4.62)$$

where \mathbf{U}_σ and \mathbf{V}_σ are $\binom{N_{ap}-1}{2} \times \binom{N_{ap}-1}{2}$ and $d \times d$ orthogonal matrices, respectively. $\boldsymbol{\Sigma}_\sigma$ is a $\binom{N_{ap}-1}{2} \times d$ diagonal matrix with r non-zero diagonal entries (the so-called *singular values* of \mathbf{C}_{oc}), where $r = \text{rank}(\mathbf{C}_{oc}) = d - 2$.

The Moore-Penrose pseudo-inverse is then given by (Bretscher, 2001):

$$\mathbf{C}_{oc}^+ = \mathbf{U}_\sigma \boldsymbol{\Sigma}_\sigma^+ \mathbf{V}_\sigma^T \quad (4.63)$$

where $\boldsymbol{\Sigma}_\sigma^+$ is a diagonal matrix whose r non-zero diagonal entries are the reciprocals of the corresponding non-zero entries in $\boldsymbol{\Sigma}_\sigma$.

In Chapter 3, we showed that the effect of wrapping on the RSC pseudo-inverse solution to Equation (4.45) can be reliably limited to an image shift for wrap-invariant patterns. An analogous result and proof apply to closure-based RSC imaging:

Proposition 4.4.5. *For wrap-invariant closure mappings, the error induced by wrapping of the (generalized) closure phases can be limited to an image shift using the standard Moore-Penrose pseudo-inverse estimator.*

Proof: See Appendix C.3 □

We can leverage this result to establish an analogous result pertaining to the well-known

non-linear least-squares formulation due to Gorham *et al.* (1989) discussed in Section 4.2. Namely we have the following Corollary:

Corollary 4.4.6. *Given a wrap-invariant closure mapping, the set of solutions minimizing Objective Ψ_2 in Equation 4.3 differ from the true solution by an image shift in the noiseless case.*

Proof: For Objective Ψ_2 to be zero, clearly each term (i.e. each squared-residual) in the summation must be zero. Hence we must have $e^{j\hat{\beta}_{i,cl}} = e^{j(\theta_{i1}+\theta_{i2}+\theta_{i3})} \mid \forall i$. These constraints are clearly equivalent to the linear phase constraints solved by the family of solutions in Equation (4.61). Hence the solution sets are the same in the noiseless case. \square

The error covariance of elements in the pseudo-inverse solution can be expressed as (Montgomery *et al.*, 2006):

$$\Sigma_{pinv} = \hat{\mathbf{V}}_{\sigma}^T \Sigma_{+}^2 \hat{\mathbf{V}}_{\sigma} \quad (4.64)$$

where $\hat{\mathbf{V}}_{\sigma}$ is obtained by omitting the final 2 columns of \mathbf{V}_{σ} which form a basis for the nullspace of \mathbf{C}_{oc} .

4.4.2 Selection of the N-Spectra

In this section, we describe a strategy for obtaining a near-optimal linearly-independent set of generalized closure phases. This strategy is founded on the notion of *minimum cycle basis* from graph theory.

Definition 4.4.4. (Minimum cycle basis of a directed graph): Let each edge of the graph be assigned a positive weight c_k , and the weight of a cycle be defined as the sum of the weights of its constituent cycles.

With these Definitions, we can then ask for the set of linearly-independent oriented cycles which spans the cycle space and has minimum total weight. We call such a set the *minimum cycle basis* of a directed graph G .

We next state a fundamental lemma that establishes a practical method for obtaining a minimum cycle basis, for which the proof can be found in many textbooks on graph theory

(see e.g. Gross and Yellen (2006)).

Lemma 4.4.7. (*Optimality of the Greedy Algorithm*): Consider a vector space V , a set of vectors which span V , and a corresponding set of weights for each of these vectors. A basis of minimum total weight can be found by first sorting the vectors in decreasing order of weight, and then selecting vectors for the basis in this order if and only if they are linearly-independent of previous selections. This is known as the greedy algorithm.

From Lemma 4.4.7, we can find a set of linearly-independent closure relations of minimum total variance by employing the greedy algorithm on the set S of all possible closure relations. However, this set is equivalent to the set of all possible cyclic permutations, which becomes enormous for arrays of large size. Fortunately, in an extension of a result due to Horton (1987) for undirected graphs, Liebchen and Rizzi (2005) showed that the elements of the minimum cycle basis of a directed graph could be found among a special (and much smaller) subset of S known as the *Horton cycles*. We describe these cycles through the following definitions:

Definition 4.4.5. (*Shortest path tree*): The shortest path tree of a graph G with respect to a node A is the spanning tree which connects node A to each other node in G via a path of minimum weight.

Figure 4.3 gives a simple example of a shortest path tree for an interferometric graph in which the top-most node is the root. The construction of shortest path trees is a well-studied problem in graph theory (see e.g. Gross and Yellen (2006)) which is typically solved using a shortest-path routine such as Dijkstra's Algorithm (Dijkstra, 1959).

Definition 4.4.6. (*Horton cycle*): Given the shortest path tree T originating from an arbitrary node A , a Horton cycle is a cycle formed by connecting the endpoints of any two branches of T .

Lemma 4.4.8. (*Minimum-cycle-basis elements are Horton cycles*): All elements of the minimum-cycle-basis are Horton cycles. Therefore in searching for the minimum-cycle-basis, it suffices to search

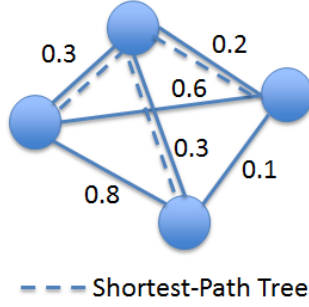


Figure 4.3: A shortest path tree

the Horton cycles. *Proof:* We summarize the Proof due to Liebchen and Rizzi (2005) in Appendix C.4. □

To exploit the Lemma above, we seek a decomposition of the weight of a cycle into the weights of its constituent edges (i.e. baselines). As the following Propositions show, at the extreme low- and high-SNR limits, the n-spectrum variance does indeed allow such a decoupling.

Proposition 4.4.9. (Low-SNR Decoupling Approximation): *In the low-SNR limit, the logarithm of the total variance of any cycle is given (up to a global additive constant) by the sum of the costs assigned to each baseline in the cycle, where the cost of a baseline k with visibility γ_k is given by:*

$$c_k = \log \frac{2}{n} - 2 \log \gamma_k.$$

Proof: This can be seen by re-writing the low-SNR approximation in Equation (4.26) as:

$$\log \sigma_{\theta_{cl}, low}^2 \approx \sum_{i=1}^o \log \frac{2}{n} - 2 \log \gamma_i + C \quad (4.65)$$

where $C = \log N_f$ is a constant with respect to both visibility and flux and therefore irrelevant for the purposes of closure selection. □

Proposition 4.4.10. (High-SNR Decoupling Approximation): *In the high-SNR limit, the total variance of any cycle is given (up to a global scaling factor) by the sum of the costs assigned to each baseline in the cycle, where the cost of a baseline i with visibility γ_i is given by: $c_i = \frac{1}{\gamma_i^2}$.*

Proof: This follows directly from the high-SNR approximation in Equation (4.25). □

Based on the Propositions above, we now have a way to construct a set of linearly-independent generalized closures of minimum variance at the low- and high- SNR extremes. As demonstrated in the next section, even at moderate flux-levels where these approximations lack precision, this closure selection method can yield significant improvements over traditional closure imaging restricted to baseline triangles. To estimate the scaling of the computational cost of the method, recall that the greedy algorithm is run on the set of Horton cycles. Since there are N_{ap} shortest-path trees, each rooted at a distinct aperture, and $m - (N_{ap} - 1)$ cycles in each of these trees, the total number of Horton cycles is $O(mN_{ap})$. Since $m = O(N_{ap}^2)$, we see that the algorithm is $O(N_{ap}^3)$, which is the same complexity as if the greedy algorithm were run on an exhaustive listing of all $\binom{N_{ap}}{3}$ three-baseline closures.

4.5 A Practical Algorithm for RSC Closure Imaging

In this section we synthesize the techniques developed in previous sections into an algorithm for RSC imaging using generalized closures. Recall that closure selection relies on estimates of the baseline visibilities $\{\gamma_i\}$. These estimates can be derived for the pairwise case by solving Equation (4.12) for γ to yield:

$$\gamma_i = \frac{\sqrt{\langle zz^* \rangle} - 2\hat{n}}{\hat{n}} \quad (4.66)$$

Similarly, for the Fizeau case, we can solve Equation (4.40) to obtain:

$$\gamma_i = \frac{\sqrt{\langle zz^* \rangle} - N_{ap}n}{n} \quad (4.67)$$

Having estimated both Fourier phases and visibilities, we can form estimates for the complex visibilities \hat{v} of the object. To estimate the image coefficients \hat{x} from these complex visibilities, we must apply a deconvolution algorithm. For this purpose we have chosen the standard sparse-recovery regularization known as Total Variation Minimization (TV-Min), whose success in reconstructing piecewise smooth objects is well-known (Becker *et al.*, 2011). The TV-Min regularization is given by:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{\alpha}} \|\mathbf{\alpha}\|_{TV} \text{ subject to: } \|\mathbf{F}\mathbf{\alpha} - \hat{\mathbf{v}}\|_2 \leq \epsilon \quad (4.68)$$

where \mathbf{F} is a partial Fourier matrix whose rows are vectorized representations of the 2D sinusoids associated with the array's measured spatial frequencies.

Algorithm 3 provides the complete set of steps we have discussed.

Algorithm 3 RSC Image Reconstruction Algorithm

1. **Estimate Visibilities** γ_i (c.f. Equations (4.66)-(4.67))
 - 2 **Select N-Spectra Relations via Minimum Cycle Basis**
 - 2.1 choose decoupling approximation according to SNR regime (c.f. Proposition 4.4.9 or 4.4.10)
 - 2.2 assign baseline edge weights accordingly
 - 2.3 enumerate the Horton cycles of the interferometric graph
 - 2.4 execute greedy algorithm on set of Horton cycles to find minimum-variance set of $m - (N_{ap} - 1)$ n-spectra
 3. **Average the Selected N-Spectra**
 4. **Unwrap the generalized closure phases corresponding to these N-Spectra** (c.f. Section 4.4.1)
 5. **Solve generalized least-squares problem with unwrapped generalized closures**
 6. **De-convolve the estimated complex visibilities to produce image reconstruction using the regularization in Equation (4.68) or other regularization technique**
-

We used the software package known as NESTA (Becker *et al.*, 2011) to execute regularization in Equation (4.68).

4.6 Algorithm Performance

In this Section, we present the results of application of Algorithm 3 to a simulated scenario of imaging a structured object space.

4.6.1 Sensitivity Limits

A useful benchmark for our results is provided by the Cramér-Rao Lower Bound (CRLB) for interferometric phase estimation in the absence of atmospheric turbulence. We begin by defining our fringe measurement model in the standard way:

$$\mathbf{p} = \tilde{\mathbf{A}} \begin{bmatrix} \mathbf{x}_{cv} \\ \mathbf{y}_{cv} \end{bmatrix} + \frac{nN_{ap}}{N_{fp}^2} \mathbf{1}_{N_{fp}^2}^T \quad (4.69)$$

where \mathbf{p} is the vector of pixel counts, \mathbf{x}_{cv} , and \mathbf{y}_{cv} are the $\binom{N_{ap}}{2}$ real and imaginary components of the complex-visibility components, respectively, $\tilde{\mathbf{A}}$ is the matrix mapping the real and imaginary parts of the complex visibilities to pixels on the focal plane (often called the visibility-to-pixel matrix *V2PM* in the literature), and $\mathbf{1}_{N_{fp}^2}^T$ is the all-ones vector of length N_{fp}^2 . Recall that in Section 1.4 we formed this mapping as a matrix \mathbf{A} , whose rows were the sinusoidal functions associated with each fringe generated by the beam combiner. A comparison between the definitions of the two mappings yields the following relation:

$$\tilde{\mathbf{A}} = \frac{2n}{N_{fp}^2} \mathbf{A} \mathbf{P}_A \quad (4.70)$$

where \mathbf{P}_A performs a permutation of the columns of \mathbf{A} to respect the ordering of the quadrature components \mathbf{x}_{cv} , and \mathbf{y}_{cv} in Equation (4.69).

Leveraging the analysis in Chapter 1 (see Section 1.4)², we can compute the Fisher-Information Matrix (FIM) for complex-visibility estimation with a Fizeau beam-combiner as:

$$\mathbf{I}_{\mathbf{x}_{cv}, \mathbf{y}_{cv}} = \left(\tilde{\mathbf{A}}^T \Lambda_{fp}^{-1} \tilde{\mathbf{A}} \right) \quad (4.71)$$

where Λ_{fp} is a diagonal matrix whose diagonal entries are the expected values of the photon counts at each detector in the array.

The CRLB for the covariance of these vector parameters is then given by:

$$\mathbf{C}_{\mathbf{x}_{cv}, \mathbf{y}_{cv}} - \mathbf{I}_{\mathbf{x}_{cv}, \mathbf{y}_{cv}}^{-1} \geq \mathbf{0} \quad (4.72)$$

where the notation \geq in this context means that the matrix difference on the left-hand side is positive-semidefinite.

²The interested reader is directed to Zmuidzinis (2003) for the original derivation

The above bound applies to single-frame fringe phasor estimation for a non-redundant array. Since we are instead evaluating phase estimation schemes for a redundant array given multiple frames of data, we must incorporate these transformations into the bound (Kay, 1993). Namely, for a redundant array, the parameter vector must be shortened to the d distinct complex-visibility phasors $\mathbf{x}_{cv,d} + (1j)\mathbf{y}_{cv,d}$. Let us define the vector $\mathbf{z}_{cv,d}$ as the concatenation of $\mathbf{x}_{cv,d}$ and $\mathbf{y}_{cv,d}$. The fringe model is then given by: $\mathbf{p} = \tilde{\mathbf{A}}\mathbf{R}\mathbf{z}_{cv,d} + \frac{nN_{ap}}{N_{fp}^2}\mathbf{1}_{N_{fp}}^T$, where \mathbf{R} is a $2\binom{N_{ap}}{2} \times 2d$ matrix mapping the real and imaginary parts of the distinct complex visibilities to those of the $2\binom{N_{ap}}{2}$ generated fringes. The FIM then becomes:

$$\mathbf{I}_{\mathbf{z}_{cv,d}} = \left(\mathbf{R}^T \tilde{\mathbf{A}}^T \Lambda_{fp}^{-1} \tilde{\mathbf{A}} \mathbf{R} \right) \quad (4.73)$$

Defining the function $g_\theta = \arctan \frac{y_{cv,d}}{x_{cv,d}}$, the CRLB for the phase covariance matrix Σ_θ can be expressed as:

$$\Sigma_\theta - \frac{1}{N_{frames}} \frac{\partial \mathbf{g}_\theta(\mathbf{z}_{cv,d})}{\partial \mathbf{z}_{cv,d}} \mathbf{I}_{\mathbf{z}_{cv,d}}^{-1} \frac{\partial \mathbf{g}_\theta(\mathbf{z}_{cv,d})^T}{\partial \mathbf{z}_{cv,d}} \geq \mathbf{0} \quad (4.74)$$

where $\frac{\partial \mathbf{g}_\theta(\mathbf{z}_{cv,d})}{\partial \mathbf{z}_{cv,d}}$ is the Jacobian matrix of the distinct Fourier phases with respect to the distinct complex-visibility phasors.

Since the left-hand-side of Inequality (4.74) is positive semi-definite (PSD), and the diagonal entries of a PSD matrix must be non-negative, we arrive at the following useful bound for the variance of the estimated phases:

$$var(\theta_i) \geq \frac{1}{N_{frames}} \left[\frac{\partial \mathbf{g}_\theta(\mathbf{z}_{cv,d})}{\partial \mathbf{z}_{cv,d}} \mathbf{I}_{\mathbf{z}_{cv,d}}^{-1} \frac{\partial \mathbf{g}_\theta(\mathbf{z}_{cv,d})^T}{\partial \mathbf{z}_{cv,d}} \right]_{ii} \quad (4.75)$$

4.6.2 Simulation

In this section, we provide the results of a simulation in which Algorithm 3 was applied to both interferometric architectures. For these simulations we used an RSC pattern of the Y-pattern type as shown in Figure 4.4. The corresponding UV-sampling for this pattern is displayed in Figure 4.5. To prevent aliasing on the simulated focal plane, this aperture

pattern was mapped onto the Golay non-redundant pattern (Golay, 1970) shown in Figure 4.6 for fringe generation; the model emulates the redundant-to-non-redundant pupil re-mapping technique developed by Perrin *et al.* (2006) and described in the Introduction. Our simulation assumed Poisson-distributed shot-noise and idealized detectors with zero read noise, and so is representative of a shot-noise-dominated scenario. For bispectrum observables, we used the unbiased estimator reported by Gordon, J. A. and Buscher, D. F. (2012):

$$g_{ub} = z_a z_b z_c - |z_a|^2 - |z_b|^2 - |z_c|^2 + 2N_{ap}\tilde{n} \quad (4.76)$$

where z_a , z_b , and z_c are the fringe phasors associated with the three sides of a bispectrum triangle, and \tilde{n} is an estimate of the total number of photons per aperture incident upon the focal plane.³

The target selected was NASA's Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations (CALIPSO) satellite for which the *truth* image (Hill, 2008) is shown in Figure 4.7. Figure 4.8 shows the image at the resolution attainable by the pattern. Two flux levels (2000 photoelectrons/aperture/frame, and 500 photoelectrons/aperture/frame) were considered for these simulations. Comparative error analysis using Equation (4.60) showed a lower predicted Root-mean-squared (RMS) phase error for the decoupling approximation in Proposition 4.4.10 than that in Proposition 4.4.9, and hence the former was chosen in Algorithm 3. Bispectra and n-spectra were integrated for $5e4$ frames, which corresponds approximately to an 8-minute observation time if we assume a typical frame duration of 10 msec. An implementation of the Bellman-Ford-Moore shortest-path algorithm (O'Connor, 2012) was used to generate a minimum cycle basis as per Algorithm 3.

To show the potential impact of generalizing the closure phase notion, Algorithm 3 was compared with an analogous algorithm which instead used the minimum-variance set of $\binom{N_{ap}-1}{2}$ independent traditional (i.e. three-baseline) closure phases. This minimum set was

³We did not derive the more complicated bias corrections for higher-order nspectra as empirical results indicated that at the flux levels considered, the importance of these bias corrections declined with the order of the nspectra.

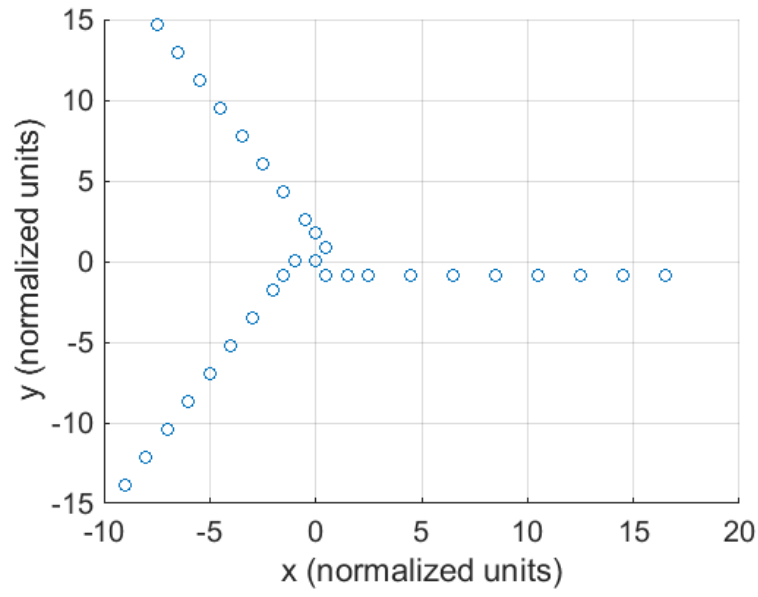


Figure 4.4: *RSC aperture pattern used in simulation*

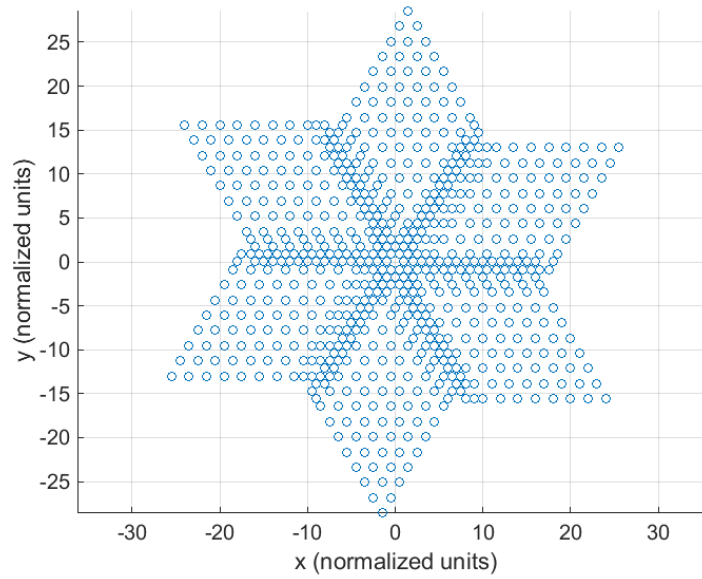


Figure 4.5: *UV-sampling for RSC pattern*

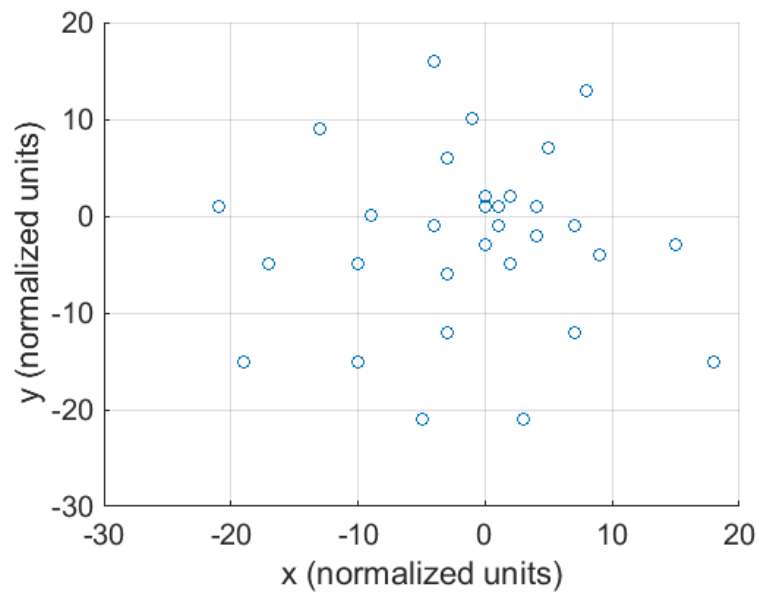


Figure 4.6: *Golay non-redundant beam-combiner pattern*

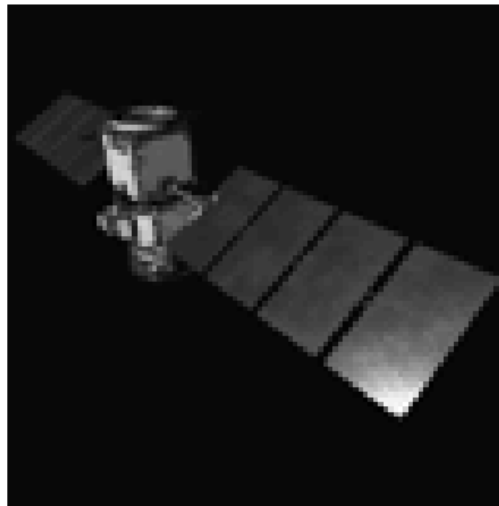


Figure 4.7: *Truth image for simulation: the CALIPSO satellite*

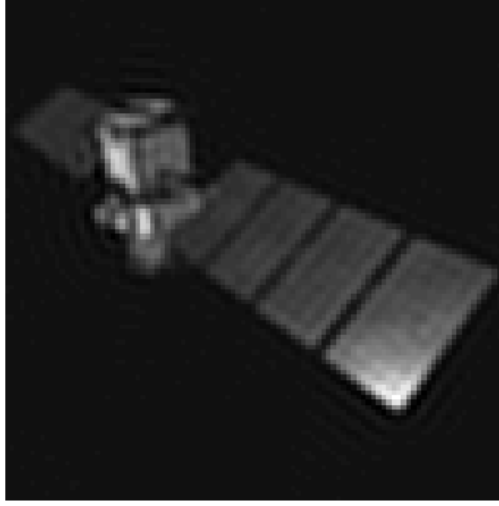


Figure 4.8: *Truth image at the resolution of the interferometric pattern*

generated via application of the greedy algorithm described above to the set of all $\binom{N_{ap}-1}{2}$.

Numerical results for the pairwise and Fizeau architectures in the higher flux scenario are shown in Figures 4.9 and 4.10, respectively. In these plots, the RMS Fourier-phase errors for 10 independent simulation trials are plotted as a function of their corresponding visibilities for the basic solution. Median-filtered versions of the predicted standard errors based on Equations (4.60) and (4.64) are also shown for both basic and pseudo-inverse solutions, respectively.⁴

As expected, the Fizeau architecture proves to be more sensitive than the Pairwise architecture, which corroborates analysis by Zmuidzinas (2003). Figures 4.11 and 4.12 respectively show Fizeau performance for the higher flux scenario at a shorter integration time ($1e3$ frames), and for the lower flux scenario at the long integration time ($5e4$ frames). Note that in the use of independent generalized closure phases according to Algorithm 3 outperforms the analogous scheme in which only traditional closures are used; the

⁴The actual predictions exhibit oscillations, which reflect the particular structure of our mapping between Fourier phases and GC's. For clarity we have created smoothed versions of the predicted curves by creating bins with edges at visibilities $\{5, 10, 20, 40, 80, 160, 320\} * 10^{-3}$, and displaying the median values in each bin.

root-mean-squared (RMS) errors are noticeably lower for the former.

It is interesting to examine performance relative to the CRLB described in Section 1.4. Recall that our CRLB is an *atmosphere-oracle* bound and hence it should bound phase estimation accuracy for the case of a stationary fringe averaged over many frames. The variance of such an estimator can be easily calculated by leveraging the analysis in Section 4.3. Namely the pseudovariance of the phasor is given by applying Equation (4.40) after averaging, so that SNR of the averaged fringe phasor $\bar{\mathbf{z}}$ is given by:

$$\sigma_{\bar{\mathbf{z}}}^2 = \frac{1}{\text{SNR}_{\bar{\mathbf{z}}}} = \frac{\sigma_{\mathbf{z}}^2}{\|E[\bar{\mathbf{z}}]\|^2} = \frac{N_{ap}}{2N_{frames}n\gamma^2} \quad (4.77)$$

We compute the predicted phase RMS for the atmosphere-oracle by taking the square-root of this expression, and add it to the plots. From the plots, we see that CRLB is virtually identical to this expression, except for isolated dips in the CRLB. These dips correspond to the redundant baselines, which as expected, admit increased phase sensitivity due to their multiplicity.

In the higher-flux scenario ($n = 2e3$), our algorithm's phase-estimation performance is within a factor of 2 from the CRLB for the vast majority of sampled visibilities. As the flux drops to $n = 5e2$, the performance begins to diverge from the CRLB as the closure phase variance leaves the regime described by Equation (4.25) and enters the regime described by Equation (4.24). That is, the variance is no longer accurately modeled simply by a linear combination of the individual phase variances of the form in Equation (4.77); rather it becomes inversely-proportional to the product of the squared-visibility. The rapid decrease in fidelity of the closure phase at low-SNR is well-known in interferometry (Kulkarni *et al.*, 1991).

Sample image reconstructions are shown in Figure 4.13. The reconstructions derived from generalized closures show greater fidelity to the true image than those derived from traditional closures, and thereby corroborate the numerical analysis presented in the plots.

It is noteworthy that while our algorithm attempts to find a minimum-variance set of n -spectra, this set is not necessarily the optimum set for phase estimation; our algorithm

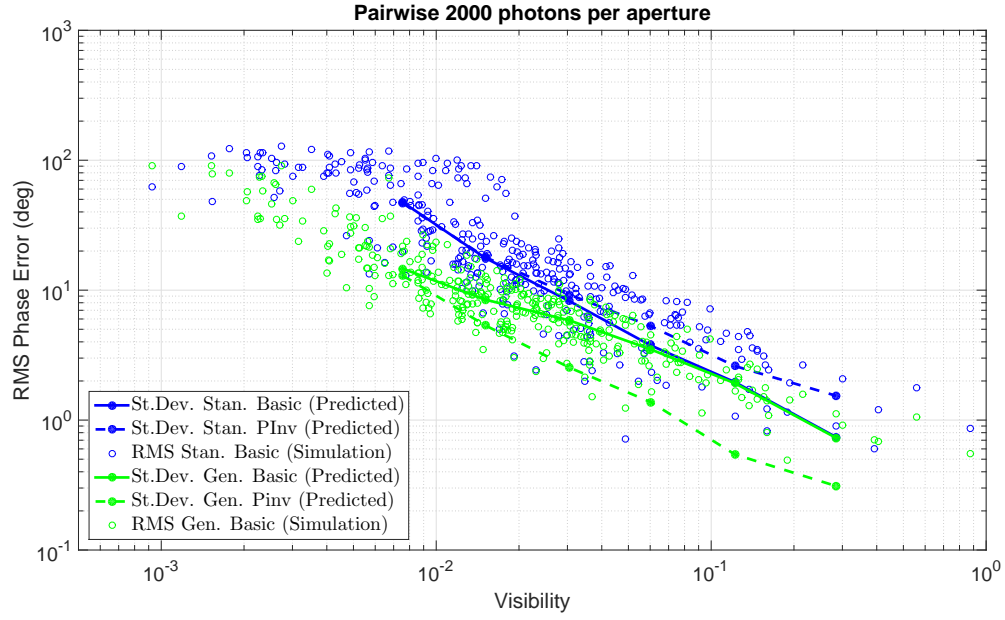


Figure 4.9: Pairwise Phase Recovery Results for Flux $n = 2e3$ pe/ap/frame, $5e4$ frames

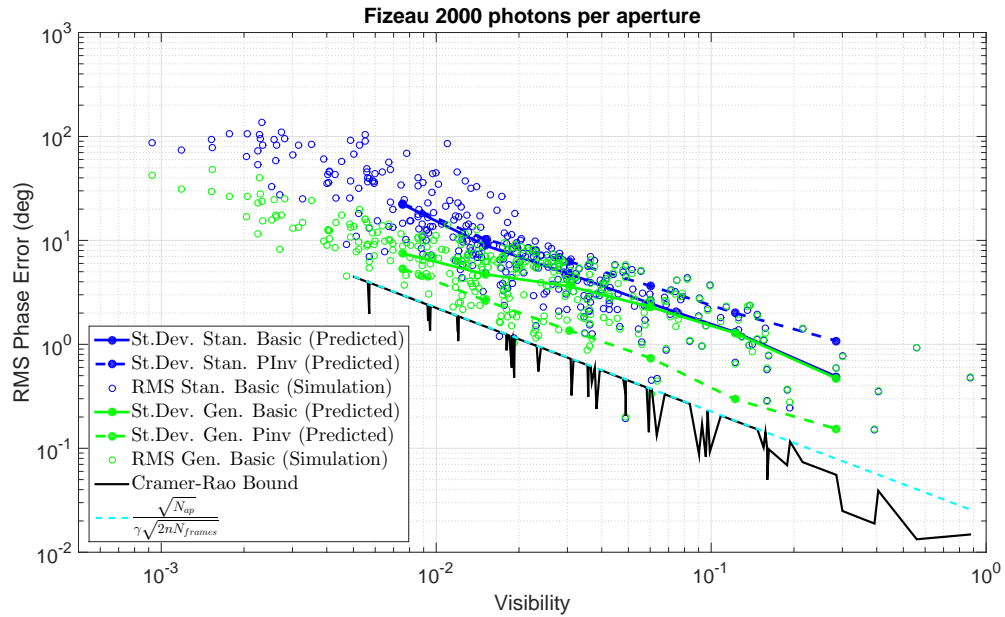


Figure 4.10: Fizeau Phase Recovery Results for Flux $n = 2e3$ pe/ap/frame, $5e4$ frames

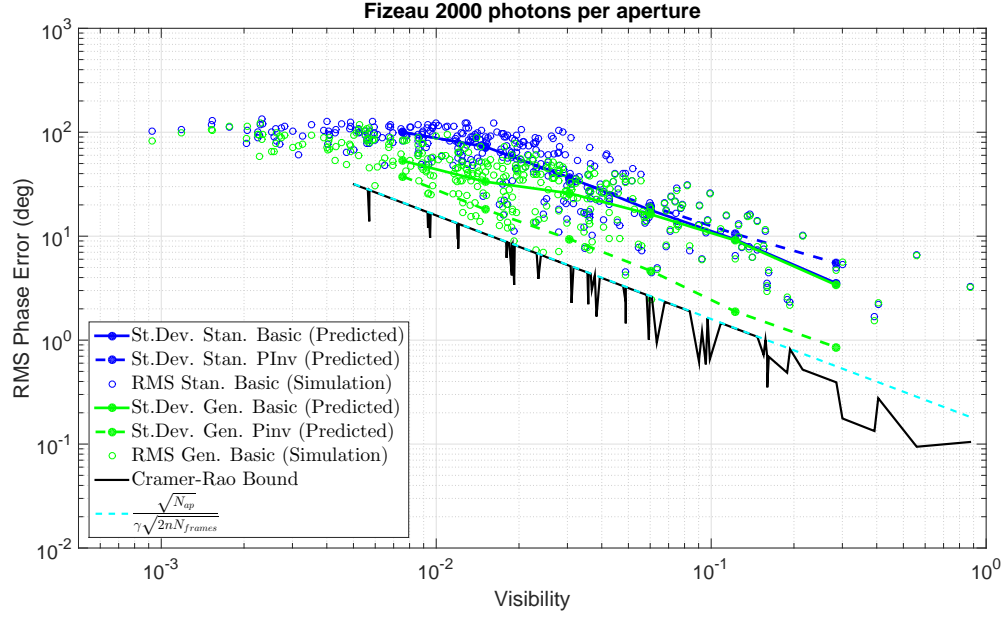


Figure 4.11: Fizeau Phase Recovery Results for Flux $n = 2e3$ pe/ap/frame, 1e3 frames

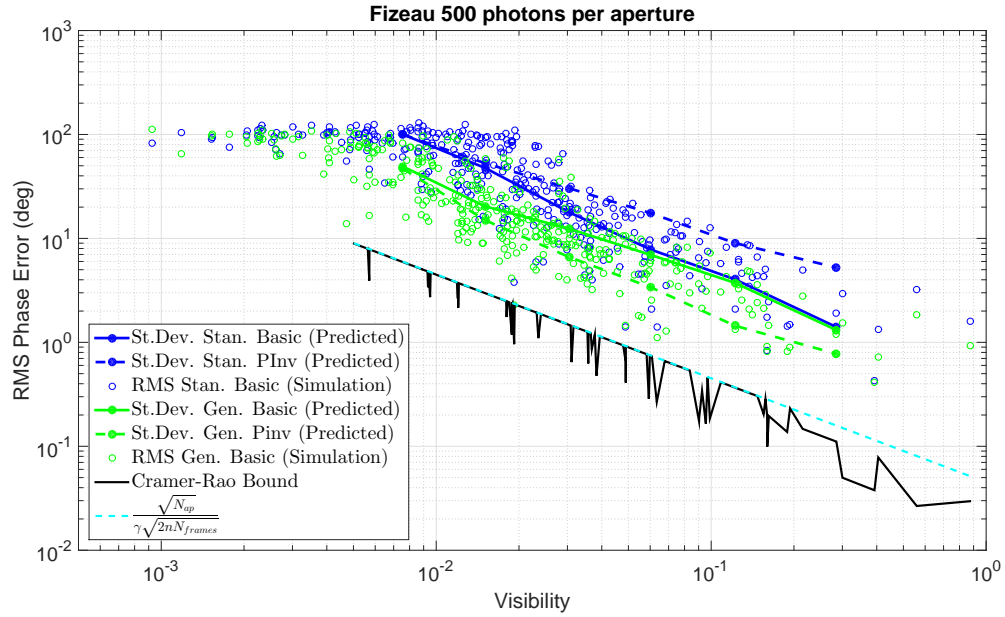


Figure 4.12: Fizeau Phase Recovery Results for Flux $n = 5e2$ pe/ap/frame, 5e4 frames

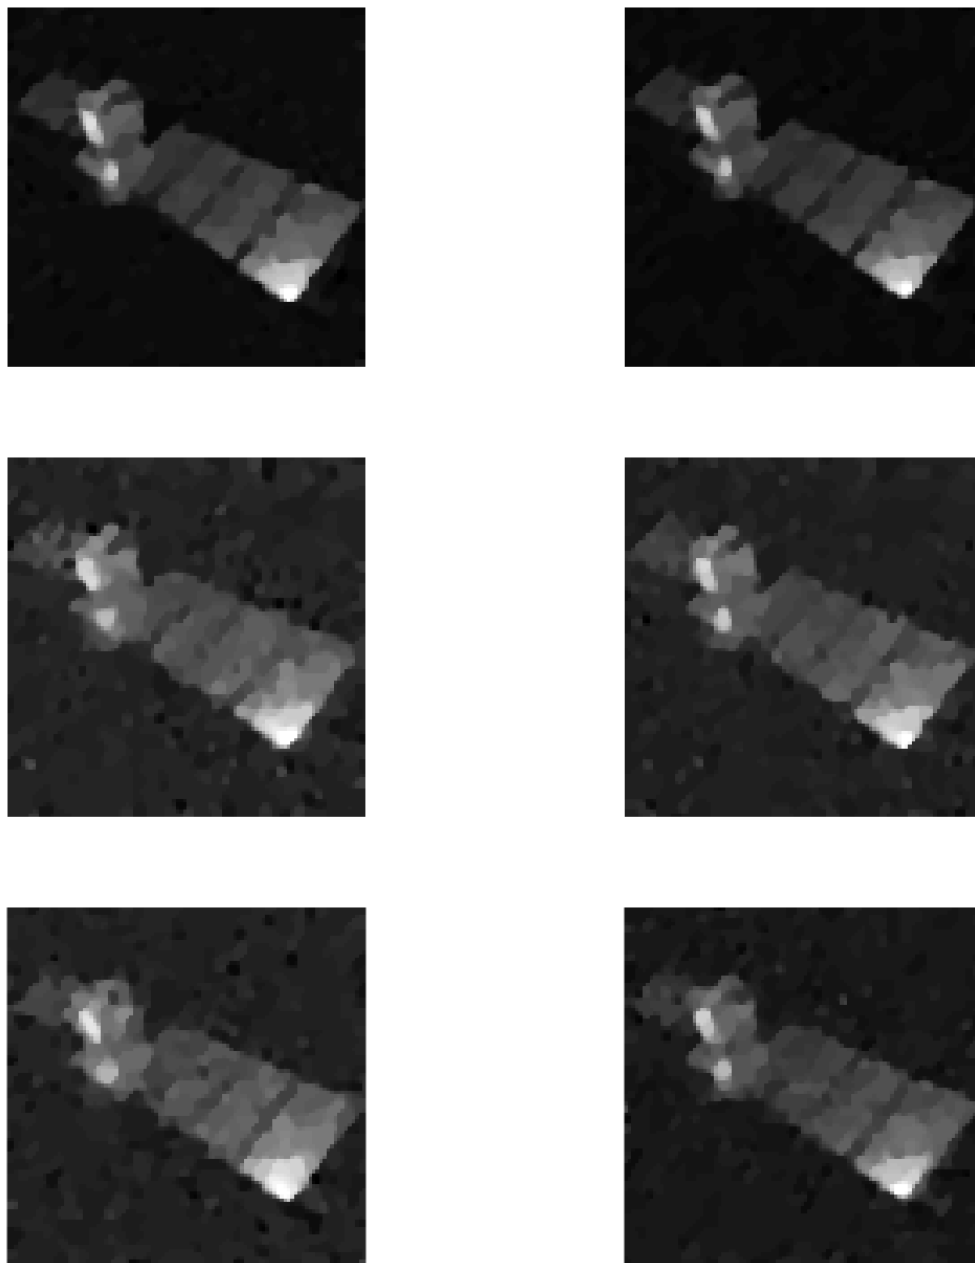


Figure 4.13: *Fizeau Image Reconstruction Results. The reconstructions in the left column used traditional, three-baseline observables, whereas those in the right column used generalized closures selected according to Algorithm 3. (top row) $n = 2e3$ photoelectrons/aperture/frame (pe/ap/frame), $5e4$ frames, (middle row) $n = 2e3$ pe/ap/frame, $2e3$ frames, (bottom row) $n = 5e2$ pe/ap/frame, $5e4$ frames.*

does not account for the conditioning of the matrix \mathbf{C}_{oc} resulting from a particular choice of nspectra. In fact there may be some cases in which the use of traditional closures results in lower phase-estimation error due to better conditioning in the corresponding \mathbf{C}_{oc} . In practice the decision to use traditional or generalized closures can be resolved by examining the phase errors predicted by Equation (4.60), for basic solutions, or Equation (4.64) for pseudo-inverse solutions.

4.6.3 Generalized Closures in Non-Linear Least Squares Approaches

This paper has proposed both novel reconstruction methodology as well as new interferometric observables. A natural question arises as to whether the advantages of the latter are retained with other, existing methodology. In particular, we consider the class of techniques which solve the non-linear least squares inference problem involving bispectrum phasors (Gorham *et al.*, 1989), (Negrete-Regagnon, 1996) (see Section 4.2).

While these algorithms bring the complexities involved in non-linear optimization (e.g. possible stagnation in local minima and/or slow convergence), they also feature the ability to utilize any number of closure relations. Algorithm 3 restricts attention to closure relations which form cycle bases, thereby guaranteeing well-posed Fourier phase recovery while keeping the size of the associated CVP-unwrapping problem tractable. However, as noted by Kulkarni *et al.* (1991), fitting to all $\binom{N_{ap}}{3}$ closures will in principle improve estimation performance when the per-frame flux is low, due to the fact that closure phases de-correlate as the flux decreases ⁵. Specifically, as computed by Kulkarni *et al.* (1991) for traditional closures in the pairwise case, the correlation coefficient μ of two closures sharing a common baseline approaches $\frac{1}{3}$ in the high-SNR regime (i.e. for $\hat{n}\gamma^2 \gg 1$) as one would expect ⁶. On the other hand, the correlation coefficient in the low-SNR regime can be accurately modeled as $\mu \approx \frac{1}{4}\gamma^4\hat{n}^2$.

To assess the phase-estimation performance provided by generalized closures relative

⁵The interested reader is directed to Buscher (2015) for a dedicated discussion of this issue

⁶Here Kulkarni *et al.* (1991) assumes a common visibility γ among all baselines in the closure

to that using the complete set of traditional closures, we consider the algorithm developed first by Gorham *et al.* (1989) which works the normalized bispectrum instead of the closure phase. We used the MATLAB[®] non-linear least-squares solver (NLS) known as *lsqnonlin* (MATLAB, 2014) to minimize the objective Ψ_2 in Equation (4.3). This solver obtains a quadratic approximation of the objective at each step and moves towards the minimum of this approximation, at which point another approximation is computed and the process repeats. The weights in Ψ_2 were set as in Gorham *et al.* (1989):

$$w_i = \frac{\hat{\sigma}_{g_i}}{\|\bar{g}_i\|} \quad (4.78)$$

where $\|\bar{g}_i\|$ is the magnitude of the averaged unbiased nspectrum, and $\hat{\sigma}_{g_i}$ is the empirical standard deviation of its quadrature components.

Results of applying this solver to the lower-SNR scenario of the previous section are given in Figure 4.14. A single randomly-chosen initialization point was used in all cases. The top row shows reconstructions and associated convergence times for the NLS algorithm using traditional closures (left) and generalized closures selected from a minimum cycle basis (middle), respectively. The bottom row gives analogous results for the case in which all $\binom{N_{ap}}{3}$ traditional closures are used. Three snapshots at iteration counts 20, 40, and 80 are shown. Iterations after 80 resulted in negligible change to the image quality. These results suggest that generalized closures can provide at least the sensitivity of traditional closures. At the same time, they benefit from the increased convergence speed afforded by a problem size proportional to the size of cycle basis (i.e. $\propto N_{ap}^2$) as opposed to the size of the set of all closures (i.e. $\propto N_{ap}^3$). These results hence show that generalized closures can serve as more efficient sources of phase information than traditional closures.

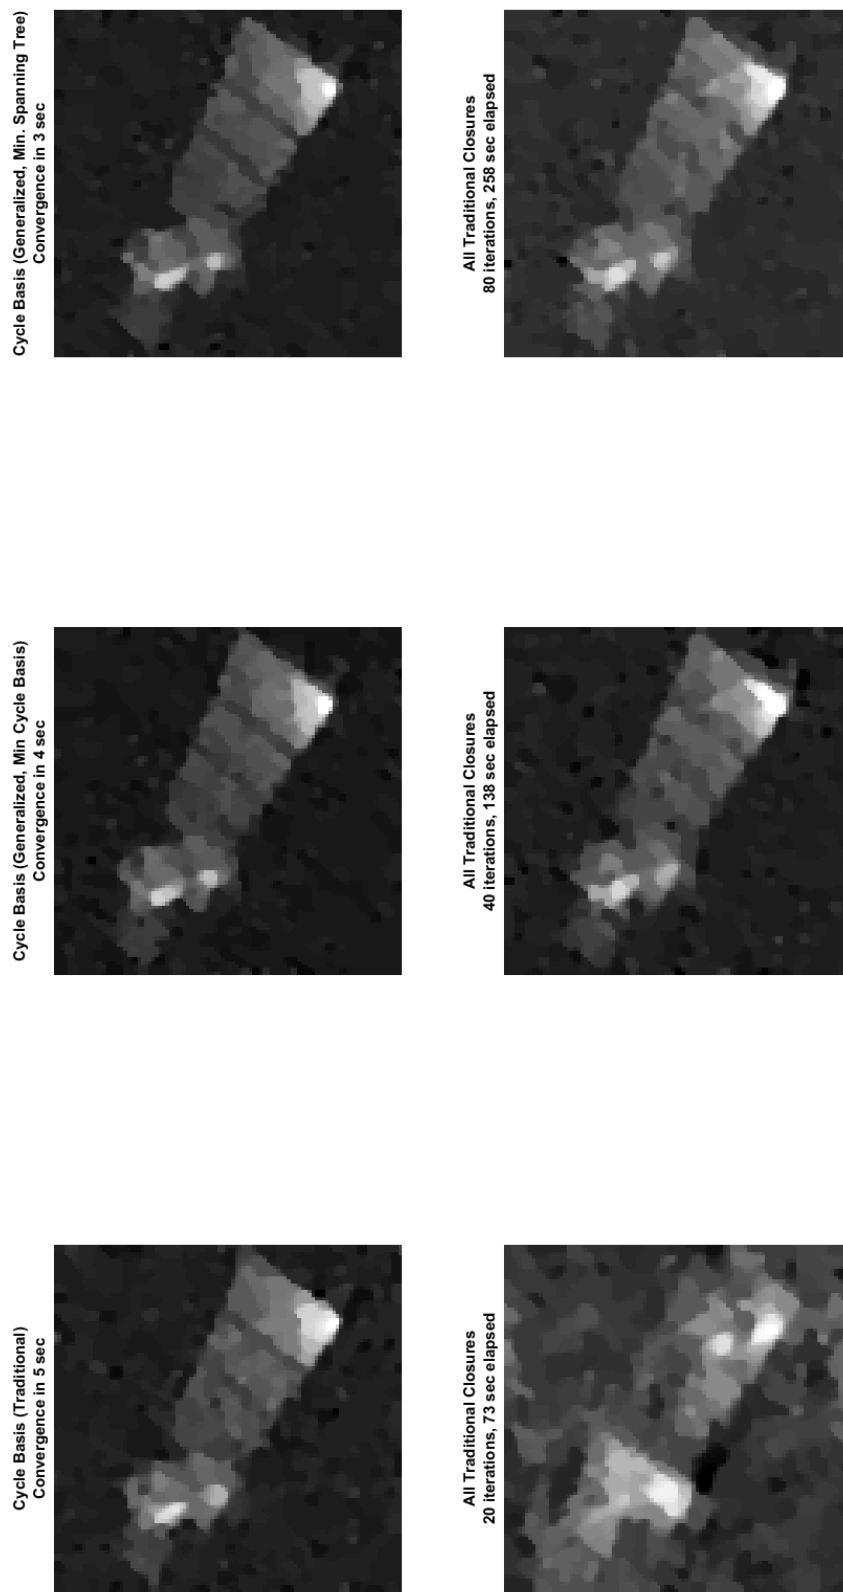
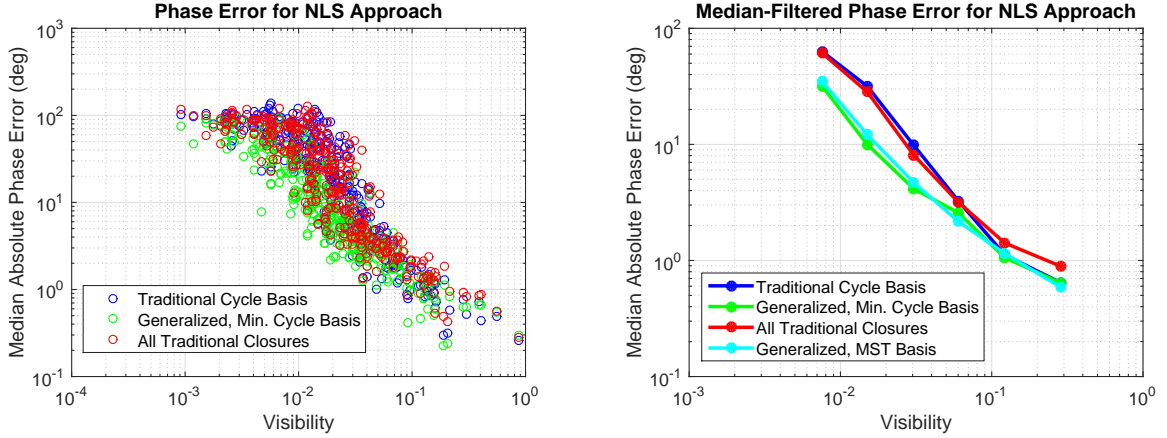


Figure 4.14: (top row) Reconstructed images and convergence times for NLS algorithm with traditional-closure basis (left), generalized-closures from minimum cycle basis (middle), and generalized-closures from minimum-variance spanning tree (right), and (bottom row) Reconstructed images and elapsed running times with all $\binom{N_{sp}}{3}$ traditional closures at iteration 20 (left), iteration 40 (middle), and iteration 80 (right) .

Figure 4.15: Raw (left) and median-filtered (right) Fourier phase error using the NLS approach.



In Figure 4.15, we plot phase error as a function of visibility for the NLS approach using a cycle basis of traditional closures (blue), a cycle basis of generalized closures (green), and using all $\binom{N_{ap}}{3}$ traditional closures (red). The median absolute error for 20 trials is shown in the scatter plot in the left panel. The same median-filter that was applied to the predictions in Figures 4.9-4.12 was then applied to this raw data to produce the results in the right panel of the Figure. It is clear that the sensitivity advantage of generalizing the closure relations is retained in the NLS approach for this example.

It is noteworthy that the selection of a minimum-variance set of generalized closures involves computational overhead in the execution of the greedy algorithm. However, it is anticipated that in many cases we may be able to obtain a reliable surrogate for this set by instead by forming a *fundamental cycle basis* (Liebchen and Rizzi, 2005) associated with minimum-variance spanning tree of the interferometric graph. Constructing this alternate basis, which was introduced in Section 3.5.3, obviates the need for execution of the greedy algorithm; it merely requires the execution of a Minimum-weight Spanning-tree algorithm (e.g. *Prim-Jarnik Algorithm* (Prim, 1957)) with each edge weight set to the reciprocal of the squared visibility associated with baseline, which is again used a proxy for the baseline's phase variance ⁷. The elements of the corresponding minimum-variance spanning tree

⁷We used a MATLAB[®] implementation of Prim's algorithm which is available online (Greenbaum, 2007)

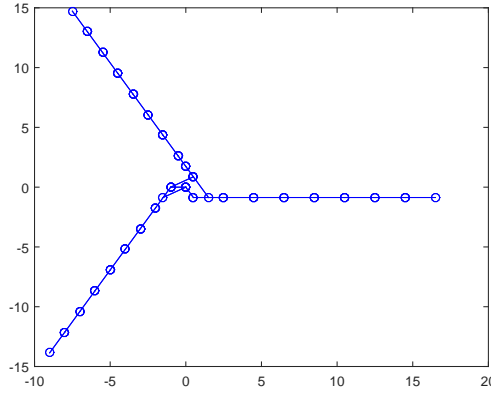


Figure 4.16: *The minimum-variance spanning tree*

(MVST) cycle basis are then found by closing the edges of this spanning tree, which is shown in Figure 4.16. Indeed, many of the Horton cycles comprising the minimum cycle basis computed for this example pattern and scenario are also members of the MVST basis. Not surprisingly, the reconstruction with this basis shown in the upper right corner of Figure 4.14 approaches of the quality that for the minimum-cycle-basis set (compare with image in the middle of the top row in Figure 4.14). Moreover the median absolute phase error for this approach, which is shown in cyan in the right panel of Figure 4.15, is very similar to that observed with the Minimum Cycle Basis.

While the analysis of Kulkarni *et al.* (1991) revealed the advantage of supplementing a cycle basis with additional traditional closures in the pairwise case, this advantage should in principle extend to the case of generalized closure phases in the Fizeau architecture. Our CRLB comparison above (i.e. Figures 4.10 and 4.11) confirms that there is limited scope for improvement in the regime $\frac{n}{N_{ap}} \approx 60$. In the regime explored in the lower-SNR scenario (see Figure 4.12), however, there is certainly a CRLB-gap which may be closed with appropriate selection of supplementary generalized closures. We leave the development of algorithms to perform such a selection efficiently for future work.

4.7 Conclusion

In this chapter, we have developed a novel method for interferometric imaging which employs RSC techniques and a generalized notion of bispectrum observable (the *n-spectrum*). We have established that phase estimation from these observables can be well-posed for valid RSC arrays. We have also provided a fast algorithm for selection of a minimum-variance set of these observables, which is based on the concept of *minimum cycle basis* in graph theory. Then, leveraging the lattice-theory problem formulation first proposed by Lannes and Anterrieu (1999) for unwrapping of the closure phases, as well as techniques from sparse recovery for image reconstruction in the presence of Fourier undersampling, we have proposed a new algorithm for image reconstruction in optical interferometry. We have shown that the performance of the phase-estimation part of our algorithm can be quantified from first principles using standard linear estimation theory. We have used simulation at different photon-flux levels to corroborate this analysis, and to show the potential advantage of performing inference using the *n-spectrum* observable as opposed to its classical counterpart - the bi-spectrum. It is our hope that both the theoretical framework employed in this chapter, as well as the practical algorithm itself, will prove useful to future users in designing RSC-based interferometric systems and processing their data.

Acknowledgements

The author would like to acknowledge Professor Jeff Fessler and his students for creating the MATLAB-based software package NUFFT (Fessler, 2003). This software was used to create the image reconstruction results shown in this chapter.

Chapter 5

Conclusions

This thesis has presented both theoretical and practical contributions to the study of robust imaging in optical interferometry. Our principal theoretical contribution is that we have developed the notion of the *wrap-invariant* RSC pattern. We have shown that these patterns enable unique recovery of the Fourier phase of the object under observation. These results are fundamental in the sense that they apply to the entire family of RSC techniques which fit the Fourier phase either directly to the complex visibilities measured by the interferometer, or atmosphere-invariant combinations of these such as the bispectrum. We summarize our uniqueness results in Table 5.1, classifying them according to the Approach type (i.e. Phase vs. Phasor), and the observable used in the phase recovery (i.e. direct measurements of the fringe phase or phasor vs. an atmosphere-invariant combination thereof). References to other work in the literature employing each observable-approach pair are also provided for context.

We have leveraged these theoretical results to construct a novel algorithmic framework for image reconstruction. In our framework, we have generalized classical atmosphere-invariant observables to improve estimation of the Fourier phase of complex objects. The ultimate goal of this effort was to achieve near-immunity to the effect of atmospheric turbulence in terms of Fourier-phase estimation accuracy, or more concretely, to show performance approaching the atmosphere-oracle Cramer Rao Lower Bound. Using our algorithm, we

provide strong theoretical and empirical evidence that we can achieve this goal in the regime in which the per-frame photon flux is on the order of 10 photons per interferometric fringe (i.e. $\frac{n}{N_{ap}}$ is on the order of 10). We have seen that the increased sensitivity of our generalized observables extends to other existing phase-recovery algorithms used in the field. Lastly, we demonstrate that the accurate phase estimation provided by our framework allows powerful techniques from Compressed Sensing to generate high-quality interferometric imagery.

Approach \ Observables	Direct Fringe Measurements Thesis Result (Ch. 3)	Related Ref.	Atmosphere-Invariant Thesis Result (Ch. 4)	Related Ref.
Phase Approach	Wrap-invariant RSC pattern \implies Least-squares solution for Fourier phases from wrapped measured phases is uniquely-recoverable up to a linear ramp (i.e. image shift)	Arnot <i>et al.</i> (1985) Greenaway (1990) Wieringa (1992) Lannes and An-terrieu (1999) Kurien <i>et al.</i> (2016)	RSC pattern with wrap-invariant generalized closure (GC) mapping \implies Least-squares solution for Fourier phases from wrapped GC's is uniquely-recoverable up to a linear ramp (i.e. image shift)	Haniff (1991) Lannes (2003) Kurien <i>et al.</i> (2016)
Phasor Approach	Wrap-invariant RSC pattern \implies global minima of least-squares objective (involving residuals of complex visibilities) differ from true solution by at most an image shift in the noiseless case	Liu <i>et al.</i> (2010) Wijnholds and Noorishad (2012) Marthi and Chengalur (2014)	RSC pattern with wrap-invariant generalized closure mapping \implies global minima of least-squares objective (involving residuals of normalized n-spectra) differ from true solution by at most an image shift in the noiseless case	Gorham <i>et al.</i> (1989) Negrete-Regagnon (1996) Thiébaud and Giovannelli (2010)

Table 5.1: Summary of new uniqueness results presented in this thesis: the implications of wrap-invariant RSC patterns

References

- AGRELL, E., ERIKSSON, T., VARDY, A. and ZEGER, K. (2002). Closest point search in lattices. *IEEE Transactions on Information Theory*, **48**, 2201–2214.
- ALI, Z. S., PARSONS, A. R., ZHENG, H., POBER, J. C., LIU, A., AGUIRRE, J. E., BRADLEY, R. F., BERNARDI, G., CARILLI, C. L., CHENG, C., DEBOER, D. R., DEXTER, M. R., GROBBELAAR, J., HORRELL, J., JACOBS, D. C., KLIMA, P., MACMAHON, D. H. E., MAREE, M., MOORE, D. F., RAZAVI, N., STEFAN, I. I., WALBRUGH, W. P. and WALKER, A. (2015). Paper-64 constraints on reionization: The 21 cm power spectrum at $z = 8.4$. *The Astrophysical Journal*, **809** (1), 61.
- ARNOT, N., ATHERTON, P., GREENAWAY, A. and NOORDAM, J. (1985). Phase closure in optical astronomy. *Traitement du Signal*, **2**, 129–136.
- BABAI, L. (1986). On lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, **6**, 1–13.
- BARTELT, H., LOHMANN, A. W. and WIRNITZER, B. (1984). Phase and amplitude recovery from bispectra. *Appl. Opt.*, **23** (18), 3121–3129.
- BECKER, S., BOBIN, J. and CANDÈS, E. (2011). NESTA: A fast and accurate first-order method for sparse recovery. *SIAM J. Imaging Sci.*, **4**, 1–39.
- BESNERAIS, G. L., LACOUR, S., MUGNIER, L. M., THIÉBAUT, E., PERRIN, G. and MEIMON, S. (2008). Advanced imaging methods for long-baseline optical interferometry. *IEEE Journal of Selected Topics in Signal Processing*, **2**, 767–780.
- BLANCHARD, P., GREENAWAY, A., ANDERTON, R. and APPLEBY, R. (1996). Phase calibration of arrays at optical and millimeter wavelengths. *J. Opt. Soc. Am. A*, **13**, 1593–1600.
- BRETSCHER, O. (2001). *Linear Algebra with Applications*. Upper Saddle River, New Jersey, USA: Prentice Hall, 2nd edn.
- BUSCHER, D. (2015). *Practical Optical Interferometry*. Cambridge CB2 8BS, U.K.: Cambridge University Press, 1st edn.
- CANDÈS, E. J., ROMBERG, J. K. and TAO, T. (2006). Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, **59** (8), 1207–1223.
- CASELLA, G. and BERGER, R. (2002). *Statistical Inference*. Pacific Grove, CA: Thomson Learning, 2nd edn.

- COHEN, H. (1993). *A course in computational algebraic number theory*, Graduate Texts in Mathematics, vol. 138. Springer-Verlag.
- DIJKSTRA, E. W. (1959). A note on two problems in connexion with graphs. *NUMERISCHE MATHEMATIK*, **1** (1), 269–271.
- DONOHO, D. (2006). Compressed sensing. *Information Theory, IEEE Transactions on*, **52** (4), 1289–1306.
- EASTWOOD, R., JOHNSON, A. and GREENAWAY, A. (2009). Calculation and correction of piston phase aberration in synthesis imaging. *J. Opt. Soc. Am. A*, **26**, 195–205.
- FESSLER, J. (2003). Nufft matlab toolbox. <http://web.eecs.umich.edu/fessler/code/>.
- GALBRAITH, S. (2012). *Mathematics of Public Key Cryptography*. New York, NY, USA: Cambridge University Press, 1st edn.
- GLINDEMANN, A. (2011). *Principles of Stellar Interferometry*, Astronomy and Astrophysics Library, New York, USA: Springer, chap. Chapter 3.3: Propagation through Interferometers. 1st edn.
- GOLAY, M. J. E. (1970). Point arrays having compact, nonredundant autocorrelations. *J. Opt. Soc. Am.*, **61**, 272–273.
- GORDON, J. A. and BUSCHER, D. F. (2012). Detection noise bias and variance in the power spectrum and bispectrum in optical interferometry. *Astronomy and Astrophysics*, **541**, A46.
- GORHAM, P., GHEZ, A., KULKARNI, S., NAKAJIMA, T., NEUGEBAUER, G., OKE, J. and PRINCE, T. (1989). Diffraction-limited imaging. iii - 30 mas closure phase imaging of six binary stars with the hale 5 m telescope. *Astronomical Journal*, **98**, 1783–1799.
- GREENAWAY, A. (1994). Prospects for alternative approaches to adaptive optics. In *Adaptive Optics for Astronomy*, NATO ASI Series. Series C: Mathematical and Physical Sciences, vol. 423, Boston, USA: Kluwer Academic Publishers.
- GREENAWAY, A. H. (1990). Self-calibrating dilute-aperture optics. *Proc. SPIE 1351, Digital Image Synthesis and Inverse Optics*, **1351**, 738–748.
- GREENBAUM, A. (2007). Prim’s algorithm. Online: [https://www.math.washington.edu/greenbau/Math 381/programs/prim.m](https://www.math.washington.edu/greenbau/Math%20381/programs/prim.m), accessed: 2015-6-01.
- GROSS, J. and YELLEN, J. (2006). *Graph Theory and its Applications*. Boca Raton, FL, USA: Chapman and Hall, 2nd edn.
- HANIFF, C. A. (1991). Least-squares fourier phase estimation from the modulo 2π bispectrum phase. *J. Opt. Soc. Am. A*, **8** (1), 134–140.
- HARMANY, Z., MARCIA, R. and WILLETT, R. (2012). This is spiral-tap: Sparse poisson intensity reconstruction algorithms;theory and practice. *Image Processing, IEEE Transactions on*, **21** (3), 1084–1096.

- HILL, E. (2008). Calipso satellite. Online: http://calipsooutreach.hamptonu.edu/images/img_gallery/Satellite/index.html, accessed: 2015-6-01.
- HÖGBOM, J. (1974). Aperture synthesis with a non-regular distribution of interferometer baselines. *Astronomy and Astrophysics Supplement*, **15**, 417–426.
- HORTON, J. (1987). A polynomial-time algorithm to find the minimum cycle basis of a regular matroid. *SIAM Journal on Computing*, **16**, 358–366.
- IRELAND, M. J. (2013). Phase errors in diffraction-limited imaging: contrast limits for sparse aperture masking. *Monthly Notices of the Royal Astronomical Society*, **433** (2), 1718–1728.
- JENNISON, R. (1958). A phase sensitive interferometer technique for the measurement of the fourier transforms of spatial brightness distributions of small angular extent. *Monthly Notices of the Royal Astronomical Society*, **118**, 276.
- KAY, S. (1993). *Fundamentals of Statistical Signal Processing, Vol. I Estimation Theory*. Upper Saddle River, NY: Prentice Hall, 1st edn.
- KULKARNI, S., PRASAD, S. and NAKAJIMA, T. (1991). Noise in optical synthesis images. ii. sensitivity of an nc2 interferometer with bispectrum imaging. *J. Opt. Soc. Am. A*, **8**, 499–510.
- KURIEN, B., TAROKH, V., ASHCOM, J., RACHLIN, Y. and SHAH, V. (2016). Resolving phase ambiguities in the calibration of redundant interferometric arrays: implications for array design. *Monthly Notices of the Royal Astronomical Society*, **(submitted)** (X), YY–YY.
- KURIEN, B. G., RACHLIN, Y., SHAH, V. N., ASHCOM, J. B. and TAROKH, V. (2014). Compressed sensing techniques for image reconstruction in optical interferometry. In *Imaging and Applied Optics 2014*, Optical Society of America, p. SM2F.3.
- LABEYRIE, A., LIPSON, S. and NISENSEN, P. (2006). *An Introduction to Optical Stellar Interferometry*. New York, NY: Cambridge University Press, 1st edn.
- LACOUR, S., THIÉBAUT, E. and PERRIN, G. (2007). High dynamic range imaging with a single-mode pupil remapping system: a self-calibration algorithm for redundant interferometric arrays. *Monthly Notices of the Royal Astronomical Society*, **374**, 832–846.
- LANNES, A. (2003). Phase closure imaging. In *Advances in Imaging and Electron Physics*, vol. 126, Boston, USA: Academic Press.
- and ANTERRIEU, E. (1999). Redundant spacing calibration: phase restoration methods. *Journal of the Optical Society of America A*, **16**, 2866–2879.
- LENSTRA, A. K., LENSTRA, H. W. and LOVÁSZ, L. (1982). Factoring polynomials with rational coefficients. *Mathematische Annalen*, **261** (4), 515–534.
- LIEBCHEN, C. and RIZZI, R. (2005). A greedy approach to compute a minimum cycle basis of a directed basis. *Inf. Process. Lett.*, **94**, 107–112.
- LIU, A., TEGMARK, M., MORRISON, S., LUTOMIRSKI, A. and ZALDARRIAGA, M. (2010). Precision calibration of radio interferometers using redundant baselines. *Monthly Notices of the Royal Astronomical Society*, **408**, 1029–1050.

- LIU, L., HE, Y., ZHANG, J., JIA, H. and MA, J. (2014). Optimum linear array for aperture synthesis imaging based on redundant spacing calibration. *Optical Engineering*, **53**, 1–8.
- MARTHI, V. R. and CHENGALUR, J. (2014). Non-linear redundancy calibration. *Monthly Notices of the Royal Astronomical Society*, **437**, 524–531.
- MARTINACHE, F. (2010). Kernel phase in fizeau interferometry. *The Astrophysical Journal*, **724**, 464–469.
- (2014). Kernel phase for interferometry with a rich aperture. In *Proceedings of Haute Provence Observatory Colloquium (23-27 September 2013)*, pp. ??–??
- MATLAB (2012). *version 8.0.0.783 (R2012b)*. Natick, Massachusetts: The MathWorks Inc.
- (2014). *version 8.4.0.150421 (R2014b)*. Natick, Massachusetts: The MathWorks Inc.
- MONTGOMERY, D., PECK, E. and VINING, G. (2006). *Introduction to Linear Regression Analysis*. Hoboken, NJ, USA: John Wiley and Sons, Inc., 4th edn.
- NEGRETE-REGAGNON, P. (1996). Phase recovery from the bispectrum aided by the error-reduction algorithm. *Opt. Lett.*, **21** (4), 275–277.
- NEWMAN, M. (1972). *Integral Matrices*. New York, NY: Academic Press, 1st edn.
- O’CONNOR, D. (2012). The bellman-ford-moore shortest path algorithm. On-line: <http://www.mathworks.com/matlabcentral/fileexchange/38129-the-bellman-ford-moore-shortest-path-algorithm>, accessed: 2015-10-01.
- PEARSON, T. J. and READHEAD, A. C. S. (1984). Image formation by self-calibration in radio astronomy. *Annual Review of Astronomy and Astrophysics*, **22** (1), 97–130.
- PERRIN, G., LACOUR, S., WOILLET, J. and THIÉBAUT, E. (2006). High dynamic range imaging by pupil single-mode filtering and remapping. *Monthly Notices of the Royal Astronomical Society*, **373**, 747–751.
- PRIM, R. C. (1957). Shortest connection networks and some generalizations. *The Bell System Technical Journal*, **36** (6), 1389–1401.
- READHEAD, A., NAKAJIMA, T., PEARSON, T., NEUGEBAUER, G., OKE, J. and SARGENT, W. (1988). Diffraction-limited imaging with ground-based optical telescopes. *Astronomical Journal*, **95**, 1278–1296.
- RODDIER, F. (1988). Passive versus active methods in optical interferometry. In F. Merkle (ed.), *European Southern Observatory Conference and Workshop Proceedings, European Southern Observatory Conference and Workshop Proceedings*, vol. 29, pp. 565–574.
- ROGSTAD, D. H. (1968). A technique for measuring visibility phase with an optical interferometer in the presence of atmospheric seeing. *Appl. Opt.*, **7** (4), 585–588.
- RUDERMAN, D. (1994). The statistics of natural images. *Network: Computation in Neural Systems*, **5**, 517–548.

- SMITH, H. (1861). On systems of linear indeterminate equations and congruences. *Phil. Trans R. Soc. Lond.*, **151**, 293–326.
- THIÉBAUT, E. (2013). Principles of image reconstruction in interferometry. *New Concepts in Imaging: Optical and Statistical Models*, **59**, 157–187.
- THIÉBAUT, E. and GIOVANNELLI, J. F. (2010). Image reconstruction in optical interferometry. *IEEE Signal Processing Magazine*, **27** (1), 97–109.
- VAN CITTERT, P. H. (1934). Die Wahrscheinliche Schwingungsverteilung in Einer von Einer Lichtquelle Direkt Oder Mittels Einer Linse Beleuchteten Ebene. *Physica*, **1**, 201–210.
- WALKUP, J. F. and GOODMAN, J. W. (1973). Limitations of fringe-parameter estimation at low light levels*. *J. Opt. Soc. Am.*, **63** (4), 399–407.
- WANG, Z., BOVIK, A. C., SHEIKH, H. R. and SIMONCELLI, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, **13** (4), 600–612.
- WIAUX, Y., JACQUES, L., PUY, G., SCAIFE, A. M. M. and VANDERGHEYNST, P. (2009). Compressed sensing imaging techniques for radio interferometry. *Monthly Notices of the Royal Astronomical Society*, **395** (3), 1733–1742.
- WIERINGA, M. (1992). An investigation of the telescope based calibration methods ‘redundancy’ and ‘self-cal’. *Experimental Astronomy*, **2**, 203–225.
- WIJNHOLDS, S. and NOORISHAD, P. (2012). Statistically optimal self-calibration of regular imaging arrays. In *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European*, pp. 1304–1308.
- WUBBEN, D., SEETHALER, D., JALDEN, J. and MATZ, G. (2011). Lattice reduction. *IEEE Signal Processing Magazine*, **28** (3), 70–91.
- ZERNIKE, F. (1938). The concept of degree of coherence and its application to optical problems. *Physica*, **5**, 785–795.
- ZHENG, H., TEGMARK, M., BUZA, V. and DILLON, J. (2014). Miteor: a scalable interferometer for precision 21 cm cosmology. *Monthly Notices of the Royal Astronomical Society*, **445**, 1084–1103.
- ZHOU, A. (2014). Clll lattice reduction algorithm. Online: <http://www.mathworks.com/matlabcentral/fileexchange/45149-clll-lattice-reduction-algorithm>, accessed: 2014-12-01.
- ZMUIDZINAS, J. (2003). Cramér-rao sensitivity limits for astronomical instruments: implications for interferometer design. *J. Opt. Soc. Am. A*, **20** (2), 218–233.

Appendix A

Appendix to Chapter 1

A.1 Sinusoidal Dependence of Field on Interferometer Focal Plane

As an aid for the following derivation, Figure A.1 shows the geometry involved in beam combination at the focal plane.

Note that we can express the vector position of pixel at coordinate $\boldsymbol{\rho}$ relative to the position of aperture j as:

$$\mathbf{x}_j = \mathbf{d} + \boldsymbol{\rho} - \mathbf{r}_j \quad (\text{A.1})$$

Moreover the magnitude of the wavevector \mathbf{k}_j is $\frac{2\pi}{\lambda}$, while the unit vector in the direction of \mathbf{k}_j is given by:

$$\hat{\mathbf{k}}_j = \frac{1}{\sqrt{\|\mathbf{d}\|^2 + \|\mathbf{r}_j\|^2}} [\mathbf{r}_j + \mathbf{d}] \quad (\text{A.2})$$

Hence the phase added to the field due to the path between the beam combiner and focal plane (at pixel $\boldsymbol{\rho}$) is given by:

$$\mathbf{k}_j \cdot \mathbf{x}_j = \frac{2\pi}{\lambda \sqrt{\|\mathbf{d}\|^2 + \|\mathbf{r}_j\|^2}} [\mathbf{r}_j + \mathbf{d}] \cdot [\mathbf{d} + \boldsymbol{\rho} - \mathbf{r}_j] \quad (\text{A.3})$$

Expanding the dot-product and noting that $\mathbf{d} \cdot \boldsymbol{\rho} = 0$, $\mathbf{d} \cdot \mathbf{r}_j = 0$, and $\mathbf{d} \cdot \mathbf{d}$ is constant

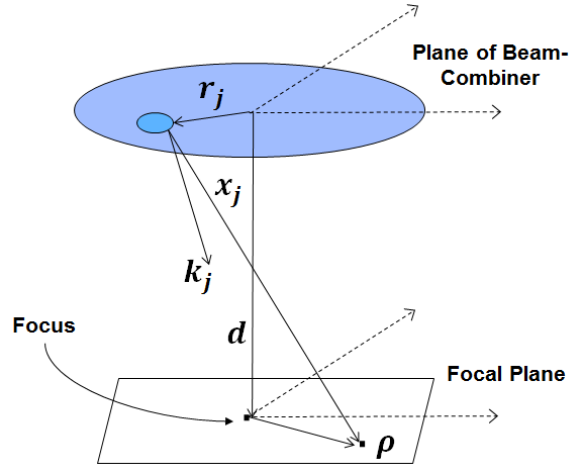


Figure A.1: *Beam-Combination Geometry*

independent of apertures, we obtain (up to a constant offset independent of aperture):

$$\mathbf{k}_j \cdot \mathbf{x}_j = \frac{2\pi}{\lambda \sqrt{\|\mathbf{d}\|^2 + \|\mathbf{r}_j\|^2}} (\mathbf{r}_j \cdot \boldsymbol{\rho}) + \phi_j \quad (\text{A.4})$$

where $\phi_j := \frac{2\pi \|\mathbf{r}_j\|^2}{\lambda \sqrt{\|\mathbf{d}\|^2 + \|\mathbf{r}_j\|^2}}$.

Appendix B

Appendix to Chapter 3

B.1 Proof of Proposition 3.5.7

In this Proposition, we propose that we can solve for the Fourier phase vector from the closure relations in two separate integer-preserving steps: the first (*Step A*) involving the SNF decomposition of $\mathbf{C}_{m \rightarrow c}$, and the second (*Step B*) involving that of \mathbf{M} . To show the validity of this two-stage decomposition, we require the following Lemma.

Lemma B.1.1. $\ker(\mathbf{C}_{m \rightarrow c}) = L$

Proof: Clearly every vector in L must be in $\ker(\mathbf{C}_{m \rightarrow c})$, since otherwise there would exist a combination of $\{\phi\}$ which would produce a non-zero closure phase. This is impossible due to the atmosphere-annihilating property of the closure phase. We must also show that L spans the kernel. Note that $\dim(\text{im}(\mathbf{C}_{m \rightarrow c})) = r = \binom{N-1}{2}$, where we use $\text{im}(\cdot)$ to denote the range of a matrix. Noting that $\mathbf{C}_{m \rightarrow c}$ has $\binom{N}{2}$ columns, we have $\dim(\ker(\mathbf{C}_{m \rightarrow c})) = \binom{N}{2} - \binom{N-1}{2} = N - 1$ by the Rank-nullity Theorem from linear algebra. Since $\dim(L) = N - 1$, the column space L indeed spans the kernel. \square

Suppose we have the SNF decomposition of $\mathbf{C}_{m \rightarrow c} = \mathbf{U}_C \mathbf{D}_C \mathbf{V}_C$. The result of Step A is then:

$$\beta_p = \mathbf{V}_C^{-1} \mathbf{D}_C^+ \mathbf{U}_C^{-1} (\mathbf{y}_{cl}^* + 2\pi \mathbf{e}_{h,cl}^*) \quad (\text{B.1})$$

Given Lemma B.1.1, Equation (B.1) can be re-written as:

$$\beta_p = \beta_0^* + \mathbf{v}_L + 2\pi \mathbf{V}_C^{-1} \mathbf{D}_C^+ \mathbf{U}_C^{-1} \mathbf{e}_{h,cl}^* \quad (\text{B.2})$$

where β_0^* is the phase measurement corresponding to the optimum least-squares solution for Equation (3.23), and \mathbf{v}_L is a vector in L . Since the elementary divisors of $\mathbf{C}_{m \rightarrow c}$ are all 1 by construction (by Lemma 2.6), the rightmost term in Equation (B.2) will be $\mathbf{0} \pmod{2\pi}$ and hence the first step is thus integer-preserving. Moreover, since \mathbf{v}_L is restricted to L , the Fourier-phase error resulting from application of the SNF-based inverse $\mathbf{V}_M^{-1} \mathbf{D}_M^+ \mathbf{U}_M^{-1}$ in Step B will be $\mathbf{0} \pmod{2\pi}$ if all elementary divisors of \mathbf{M} are 1. Hence the Proposition holds. \square

B.2 Proof of Lemma 3.7.2

In this section, we prove Lemma 3.7.2, which states the following:

Given wrap-invariance, the (column) vector $\mathbf{U}_X \mathbf{M}_K^+ \mathbf{e}_h^$ has integer entries below row 2.*

Given that we have a wrap-invariant pattern, we know that the elementary divisors of \mathbf{M} are all 1. Hence there exists an integer vector \mathbf{k}_0 such that:

$$\mathbf{e}_h^* = \mathbf{M} \mathbf{k}_0 \quad (\text{B.3})$$

Substituting Equations (C.45) and (3.26) into Equation (C.38), we obtain:

$$\mathbf{e}_\sigma = \mathbf{V}_\sigma \Sigma_\sigma^+ \mathbf{U}_\sigma^T \mathbf{U}_\sigma \Sigma_\sigma \mathbf{V}_\sigma^T \mathbf{k}_0 \quad (\text{B.4})$$

Noting that \mathbf{U}_σ is orthogonal, this equation can be simplified to

$$\mathbf{e}_\sigma = (\mathbf{V}_\sigma - \mathbf{N}) \mathbf{V}_\sigma^T \mathbf{k}_0 \quad (\text{B.5})$$

where \mathbf{N} is a matrix of the same size as \mathbf{V}_σ , which is zero except for the last three columns. These last three columns are identical to those of \mathbf{V}_σ , and hence by Lemma 4.1 comprise an orthogonal basis for the nullspace of \mathbf{M} . Noting the orthogonality of \mathbf{V}_σ , this

can be further simplified to:

$$\mathbf{e}_\sigma = \mathbf{k}_0 - \mathbf{N}\mathbf{V}_\sigma^T \mathbf{k}_0 \quad (\text{B.6})$$

To proceed, the following Definition will be useful:

Definition A.1 (The canonical basis for the nullspace of \mathbf{M}): The canonical basis $\{\mathbf{w}_i\}, i \in 1, 2, 3$ for the three-dimensional nullspace of \mathbf{M} can be derived trivially from the well-known tilt-position degeneracy in interferometry described in Section 3.5.1 (Wieringa, 1992). Namely, we can define the basis as the columns of a $(N + d) \times 3$ matrix $\mathbf{W}_{ker(\mathbf{M})}$ as follows:

$$\mathbf{W}_{ker(\mathbf{M})} = \{\mathbf{w}_1 | \mathbf{w}_2 | \mathbf{w}_3\} = \begin{bmatrix} \mathbf{0} & \Delta \mathbf{r}_x & \Delta \mathbf{r}_y \\ \mathbb{1}_{N \times 1} & \mathbf{r}_x & \mathbf{r}_y \end{bmatrix} \quad (\text{B.7})$$

where \mathbf{r}_x and \mathbf{r}_y are the x - and y -positional coordinates of the apertures associated with each row, respectively, and $\Delta \mathbf{r}_x$ and $\Delta \mathbf{r}_y$ their respective pairwise differences. \square

Note that each of the three non-zero column vectors $\{\mathbf{v}_k\}, k \in 1, 2, 3$ in \mathbf{N} can be expressed as a linear combinations of the elements of the canonical basis $\{\mathbf{w}_i\}, i \in 1, 2, 3$ defined above, i.e.

$$\mathbf{v}_k = a_1 \mathbf{w}_1 + a_2 \mathbf{w}_2 + a_3 \mathbf{w}_3 \quad (\text{B.8})$$

As in Section 3.5.1, let us again use K to denote the set of indices in \mathbf{e}_σ associated with the Fourier phases (as opposed to the piston phases), and their corresponding rows in \mathbf{N} . Hence we have:

$$\mathbf{U}_\chi \mathbf{e}_{\sigma, K} = \mathbf{U}_\chi \mathbf{k}_{0, K} - \mathbf{U}_\chi \mathbf{N}_K \mathbf{V}_\sigma^T \mathbf{k}_0 \quad (\text{B.9})$$

Since \mathbf{U}_χ is an integer matrix, the first term is clearly integral. Let us then examine the second term, and in particular, the product $\mathbf{U}_\chi \mathbf{N}_K$. By substitution from Equation (C.49), we see that for any of the three non-zero columns $\mathbf{v}_{k, K}$ of \mathbf{N}_K , we have (by simply changing the basis for the nullspace):

$$\mathbf{U}_X \mathbf{v}_{k,K} = a_1 \mathbf{U}_X \mathbf{w}_{1,K} + a_2 \mathbf{U}_X \mathbf{w}_{2,K} + a_3 \mathbf{U}_X \mathbf{w}_{3,K} \quad (\text{B.10})$$

where $\{\mathbf{w}_{j,K}\}$ are the vectors comprising the upper partition of Equation (B.7). The first term in Equation (C.51) is trivially $\mathbf{0}$ since $\mathbf{w}_{1,K}$ is the zero-vector. Now note from Theorem 3.5.2 that \mathbf{U}_X is a matrix which annihilates all the spatial frequencies in the matrix \mathbf{X} below row 2. But from Definition A.1, these spatial frequencies are identically the contents of the two columns $\mathbf{w}_{2,K}$ and $\mathbf{w}_{3,K}$ (up to a uniform scaling factor). Therefore the column vector in Equation (C.51) is zero below row 2. This means in turn that the second term in Equation (C.50) is zero below row 2, and hence that $\mathbf{U}_X \mathbf{e}_{\sigma,K}$ is integral below row 2 (since the first term is integral). \square

Appendix C

Appendix to Chapter 4

C.1 Fizeau Variance Approximations

C.1.1 Variance Decomposition

In this section we justify our approximation for the variance of the n-spectrum for the Fizeau architecture. Our goal here is not to compute all of the numerous terms in the variance, but rather to present the mathematical intuition behind the approximation we have used. The general formula for the variance of an n-spectra G is:

$$Var[G] = \left\langle \prod_{i=1}^o z_i z_i^* \right\rangle - \left\langle \prod_{i=1}^o z_i \right\rangle \left\langle \prod_{i=1}^o z_i^* \right\rangle \quad (C.1)$$

For simplicity, we analyze the standard bispectrum case ($o = 3$). The analysis extends naturally to the n-spectrum. Note that by extension of Equation (4.8) the first term can be written as:

$$\left\langle \prod_{i=1}^o z_i z_i^* \right\rangle = \sum_{p_a, p_b, p_c, p_d, p_e, p_f} \langle q(p_a) q(p_b) q(p_c) q(p_d) q(p_e) q(p_f) \rangle \times e^{i\omega_1(p_d - p_a)} e^{i\omega_2(p_e - p_b)} e^{i\omega_3(p_f - p_c)} \quad (C.2)$$

Similarly, the second term in Equation (C.1) can be written as:

$$\begin{aligned}
\left\langle \prod_{i=1}^o z_i \right\rangle \left\langle \prod_{i=1}^o z_i^* \right\rangle = & \sum_{p_a, p_b, p_c, p_d, p_e, p_f} \langle q(p_a)q(p_b)q(p_c) \rangle \langle q(p_d)q(p_e)q(p_f) \rangle \\
& \times e^{i\omega_1(p_d-p_a)} e^{i\omega_2(p_e-p_b)} e^{i\omega_3(p_f-p_c)} \quad (C.3)
\end{aligned}$$

Proceeding analogously to Equation (4.10), we utilize the general formula for the moments of a Poisson distribution to perform the following decomposition of the first factor in the summand in Equation (C.2) (Kulkarni *et al.*, 1991):

$$\langle q(p_a)q(p_b)q(p_c)q(p_d)q(p_e)q(p_f) \rangle = \quad (C.4)$$

$$+ \langle q(p_a) \rangle \langle q(p_b) \rangle \langle q(p_c) \rangle \langle q(p_d) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.5)$$

$$+ \delta_{p=p_a=p_d} \langle q(p) \rangle \langle q(p_b) \rangle \langle q(p_c) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.6)$$

$$+ \delta_{p=p_b=p_e} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_c) \rangle \langle q(p_d) \rangle \langle q(p_f) \rangle \quad (C.7)$$

$$+ \delta_{p=p_c=p_f} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_b) \rangle \langle q(p_d) \rangle \langle q(p_e) \rangle \quad (C.8)$$

$$+ \delta_{p=p_a=p_d} \langle q(p) \rangle \delta_{p'=p_b=p_e} \langle q(p') \rangle \langle q(p_c) \rangle \langle q(p_f) \rangle \quad (C.9)$$

$$+ \delta_{p=p_a=p_d} \langle q(p) \rangle \delta_{p'=p_c=p_f} \langle q(p') \rangle \langle q(p_b) \rangle \langle q(p_e) \rangle \quad (C.10)$$

$$+ \delta_{p=p_b=p_e} \langle q(p) \rangle \delta_{p'=p_c=p_f} \langle q(p') \rangle \langle q(p_a) \rangle \langle q(p_c) \rangle \quad (C.11)$$

$$+ \delta_{p=p_a=p_d} \langle q(p) \rangle \delta_{p'=p_b=p_e} \langle q(p') \rangle \delta_{p''=p_c=p_f} \langle q(p'') \rangle \quad (C.12)$$

$$+ \delta_{p=p_a=p_e} \langle q(p) \rangle \langle q(p_b) \rangle \langle q(p_c) \rangle \langle q(p_d) \rangle \langle q(p_f) \rangle \quad (C.13)$$

$$+ \delta_{p=p_a=p_f} \langle q(p) \rangle \langle q(p_b) \rangle \langle q(p_c) \rangle \langle q(p_d) \rangle \langle q(p_e) \rangle \quad (C.14)$$

$$+ \delta_{p=p_b=p_d} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_c) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.15)$$

$$+ \delta_{p=p_b=p_f} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_c) \rangle \langle q(p_d) \rangle \langle q(p_e) \rangle \quad (C.16)$$

$$+ \delta_{p=p_c=p_d} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_b) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.17)$$

$$+ \delta_{p=p_c=p_e} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_b) \rangle \langle q(p_d) \rangle \langle q(p_f) \rangle \quad (C.18)$$

$$+ \delta_{p=p_a=p_b} \langle q(p) \rangle \langle q(p_c) \rangle \langle q(p_d) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.19)$$

$$+ \delta_{p=p_a=p_c} \langle q(p) \rangle \langle q(p_b) \rangle \langle q(p_d) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.20)$$

$$+ \delta_{p=p_b=p_c} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_d) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.21)$$

$$+ \delta_{p=p_d=p_e} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_b) \rangle \langle q(p_c) \rangle \langle q(p_f) \rangle \quad (C.22)$$

$$+ \delta_{p=p_d=p_f} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_b) \rangle \langle q(p_c) \rangle \langle q(p_e) \rangle \quad (C.23)$$

$$+ \delta_{p=p_e=p_f} \langle q(p) \rangle \langle q(p_a) \rangle \langle q(p_b) \rangle \langle q(p_c) \rangle \langle q(p_d) \rangle \quad (C.24)$$

$$+ \delta_{p=p_a=p_b=p_c} \langle q(p) \rangle \langle q(p_d) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.25)$$

$$\dots + \text{other order-3 terms} \quad (C.26)$$

$$+ \delta_{p=p_a=p_b=p_c=p_d} \langle q(p) \rangle \langle q(p_e) \rangle \langle q(p_f) \rangle \quad (C.27)$$

$$130 + \text{other order-4 terms} \quad (C.28)$$

$$\dots + \text{order-6 terms} \quad (C.29)$$

The total number of partitions (203) in the summation above is given by the 6th Bell number. Below we categorize the terms by their order, which in this case we define as the largest number of common pixels in each partition. We will see that we can approximate the pseudo-variance to reasonable accuracy by including a particular subset of the Order-2 terms. Note that all terms in the analogous decomposition of Equation (C.3) are also in Equation (C.2) and hence cancel them in Equation (C.1).

C.1.2 Order-2 Terms

These terms have a pair of common pixels. We distinguish between the following two cases:

Case 1: conjugate pairs These terms are given in (C6)-(C12). The common pixels belong to the conjugate pairs, i.e. $p := p_a = p_d$, $p := p_b = p_e$, and $p := p_c = p_f$, respectively, which also appear in the pairwise case. Let us assume that the visibilities associated with ω_1 , ω_2 , and ω_3 are all roughly equal, i.e. $\bar{\gamma} = \gamma_1 = \gamma_2 = \gamma_3$. Consider the sum of the first of these terms (C6) over the focal plane, i.e.:

$$\begin{aligned}
& \sum_{p:=p_a=p_d, p_b, p_c, p_e, p_f} \langle q(p) \rangle \langle q(p_b) \rangle e^{i\omega_2 p_b} \langle q(p_c) \rangle e^{i(\omega_1 + \omega_2) p_c} \\
& \quad \times \langle q(p_e) \rangle e^{-i\omega_2 p_e} \langle q(p_f) \rangle e^{-i(\omega_1 + \omega_2) p_f} \\
& = \sum_{p:=p_a=p_d} \Lambda_p \sum_{p_b} \Lambda_{p_b} e^{i\omega_2 p_b} \sum_{p_c} \Lambda_{p_c} e^{i(\omega_1 + \omega_2) p_c} \\
& \quad \sum_{p_e} \Lambda_{p_e} e^{i\omega_2 p_e} \sum_{p_f} \Lambda_{p_f} e^{-i(\omega_1 + \omega_2) p_f}
\end{aligned} \tag{C.30}$$

The magnitude of this expression is $(N_{ap} n^5 \gamma_2^2 \gamma_3^3)$. In the case where all visibilities are equal to $\bar{\gamma}$, this expression becomes $N_{ap} n^5 \bar{\gamma}^4$. Terms (C7)-(C8) follow this same form.

Now consider term (C9):

$$\begin{aligned}
& \sum_{p=p_a=p_d, p'=p_b=p_e, p_c, p_f} \langle q(p) \rangle \langle q(p') \rangle \langle q(p_e) \rangle e^{-i\omega_2 p_e} \\
& \quad \times \langle q(p_f) \rangle e^{-i(\omega_1+\omega_2)p_f} \\
& = \sum_{p:=p_a=p_d} \Lambda_p \sum_{p':=p_b=p_e} \Lambda'_p \sum_{p_c} \Lambda_{p_c} e^{i(\omega_1+\omega_2)p_f} \\
& \quad \times \sum_{p_f} \Lambda_{p_f} e^{-i(\omega_1+\omega_2)p_f}
\end{aligned} \tag{C.31}$$

The magnitude of this expression is $N_{ap}^2 n^4 \gamma_3^2$, which, under the assumption of equal visibilities, becomes $N_{ap}^2 n^4 \bar{\gamma}^2$. Terms (C10)-(C11) follow this same form.

Case 2: mixed non-conjugate pairs

We now consider the terms in which a pair includes two distinct baselines (e.g. (C13)-(C18)). The common pixels in this case are such that one of the corresponding baselines is conjugated and the other is not. Taking, for example, the case (C13) in which $p := p_a = p_e$, we have:

$$\begin{aligned}
& \sum_{p=p_a=p_e, p_b, p_c, p_d, p_f} \langle q(p) \rangle e^{i(\omega_1-\omega_2)p} \langle q(p_b) \rangle e^{i\omega_2 p_b} \\
& \quad \times \langle q(p_c) \rangle e^{i(\omega_1+\omega_2)p_c} \langle q(p_d) \rangle e^{-i\omega_2 p_d} \langle q(p_f) \rangle e^{-i(\omega_1+\omega_2)p_f} \\
& = \sum_{p:=p_a=p_e} \Lambda_p e^{i(\omega_1-\omega_2)p} \sum_{p_b} \Lambda_{p_b} e^{i\omega_2 p_b} \sum_{p_c} \Lambda_{p_c} e^{i(\omega_1+\omega_2)p_c} \\
& \quad \times \sum_{p_d} \Lambda_{p_d} e^{-i\omega_2 p_d} \sum_{p_f} \Lambda_{p_f} e^{-i(\omega_1+\omega_2)p_f} \tag{C.32}
\end{aligned}$$

where $\gamma_{\omega_1-\omega_2}$ is the visibility of the fringe of spatial frequency $\omega_1 - \omega_2$. Note that the first factor in the Equation above clearly vanishes if such a fringe is not created by the beam combiner. Otherwise, assuming equal strength of the visibilities, the magnitude of this expression is given by $(n^5 \gamma_{\omega_1-\omega_2} \gamma_1 \gamma_2 \gamma_3^2)$, or $n^5 \bar{\gamma}^5$ in the case of equal visibilities. Comparing with any of the Case 1 terms and recalling that $\bar{\gamma} < 1$, we see that the latter will dominate

the former in the regime considered in this paper ($n \approx 1e3$, $\bar{\gamma} \ll 1$, $N_{ap} = 31$).

Case 3: terms in decomposition Equation (C.3)

This Case includes terms (C5), as well as (C19)-(C24). The common pixels are such that both or neither of the corresponding baselines are conjugated, i.e.: the cases $p_a = p_b$, $p_a = p_c$, $p_b = p_c$, $p_d = p_e$, $p_d = p_f$, $p_e = p_f$. These terms are canceled by their counterparts in Equation (C.2). \square

C.1.3 Higher Order Terms

An example of a third-order term would be the case of $p := p_a = p_b = p_d$, whose sum is given by:

$$\sum_{p:=p_a=p_b=p_d, p_c, p_e, p_f} \langle q(p) \rangle e^{i\omega_2 p} \langle q(p_c) \rangle e^{-i(\omega_1+\omega_2)p_c} \langle q(p_e) \rangle e^{-i\omega_2 p_e} \langle q(p_f) \rangle e^{-i(\omega_1+\omega_2)p_f} \quad (C.33)$$

$$= \sum_{p:=p_a=p_b=p_d} \Lambda_p e^{i\omega_2 p} \sum_{p_c} \Lambda_{p_c} e^{i(\omega_1+\omega_2)p_c} \sum_{p_e} \Lambda_{p_e} e^{-i\omega_2 p_e} \sum_{p_f} \Lambda_{p_f} e^{-i(\omega_1+\omega_2)p_f} \quad (C.34)$$

The magnitude of this expression is given by $n^4 \gamma_2^2 \gamma_3^2$. For equal visibilities, we have $n^4 \bar{\gamma}^4$. Clearly this term will be dominated by the Case 1 terms. Other high-order terms exhibit a similar sharp attenuation which allows them to be neglected for practical purposes. \square

Given the relative strengths of the terms shown above, we will retain Case 1 terms (C6)-(C12), which after summation and factoring, yield Equation (4.41). Note that this amounts to the following approximation:

$$V_{fizeau}(G) \approx \left[\prod_{i=1}^o \langle z_i z_i^* \rangle \right] - \prod_{i=1}^o \langle z_i \rangle \langle z_i^* \rangle \quad (C.35)$$

Applying analogous analysis to the Fizeau covariance yields the following approximation:

$$\sigma(G_1, G_2^*)_{fizeau} \approx \left[\prod_{k \in I \cap J} \langle z_k z_k^* \rangle - \prod_{k \in I \cap J} \langle z_k \rangle \langle z_k^* \rangle \right] \prod_{l \in I \setminus J} \langle z_l \rangle \prod_{m \in J \setminus I} \langle z_m^* \rangle \quad (C.36)$$

Note that this matches the corresponding pairwise expression in Equation (4.28). Substi-

tution then yields the final Fizeau covariance expressions in Equations (4.43)-(4.44).

C.2 Proofs of Proposition 4.4.2 and Corollary 4.4.3

C.2.1 Proof of Proposition 4.4.2

Proposition 4.4.2: For a valid RSC array, the columns $\Delta \mathbf{r}_x$ and $\Delta \mathbf{r}_y$ form a basis for the two-dimensional nullspace of \mathbf{C}_{oc} .

Proof: To see that the two columns mentioned belong to the nullspace of \mathbf{C}_{oc} , note that each solution set to the noiseless version of Equation (4.49) above remains valid after replacing each θ_{ij} with $\theta_{ij}^p = \theta_{ij} - \mathbf{z} \cdot (\mathbf{r}_i - \mathbf{r}_j)$.

We then need to establish that these two vectors span the entire nullspace of \mathbf{C}_{oc} by showing that this nullspace is two-dimensional. Suppose we have a vector \mathbf{w} which is in the nullspace \mathbf{C}_{oc} . This is equivalent to the either of the following conditions: $\mathbf{M}_\theta \mathbf{w} = \mathbf{0}$, or $\mathbf{M}_\theta \mathbf{w} \in \ker(\mathbf{C}_{mc})$. The former condition is not possible since the columns spanning the subspace K are linearly-independent. The latter condition is equivalent to the condition that $\mathbf{M}_\theta \mathbf{w} \in K \cap L$. It is well-known fact that for any two subspaces K and L , we have:

$$\dim(K \cap L) = \dim(K) + \dim(L) - \dim(K + L) \quad (\text{C.37})$$

We established in Section 3.5.1 that $\dim(K + L) = \text{rank}(\mathbf{M}) = d + N - 3$ for a valid RSC system. Also $\dim(K) = d$, since the K is spanned by d linearly-independent columns. Finally $\dim(L) = N - 1$. Substituting into Equation (C.37), we see that $\dim(K \cap L) = 2$. \square

C.2.2 Proof of Corollary 4.4.3

Corollary 4.4.3: For a valid RSC array, the mapping \mathbf{C}_{oc} is injective up to an image shift.

Proof: Proposition 4.4.2 showed that the nullspace is comprised of linear combinations of vectors $\Delta \mathbf{r}_x$ and $\Delta \mathbf{r}_y$. Note that $\Delta \mathbf{r}_x$ and $\Delta \mathbf{r}_y$ are simply scaled versions of the x - and y -spatial frequency vectors in the array, respectively. Hence adding linear combinations of these vectors to a particular RSC phase solution merely produces phase ramps in the Fourier

domain, which are equivalent to translations (or shifts) in the image domain. In other words, the mapping \mathbf{C}_{oc} is invertible up to an unknown image shift. \square

C.3 Proof of Proposition 4.4.5

We begin the Proof with the following Lemma:

Lemma C.1: The final 2 columns of \mathbf{V}_σ form a basis for the nullspace of \mathbf{C}_{oc} .

Proof: This follows from the fact \mathbf{C}_{oc} is rank-deficient by 2, and standard properties of the right singular vectors comprising \mathbf{V}_σ in the SVD. (Bretscher, 2001) \square

Note that the error resulting from application of the pseudo-inverse \mathbf{C}_{oc}^+ to the unwrapped vector of closures will be given by:

$$2\pi\mathbf{e}_\sigma = \mathbf{C}_{oc}^+(2\pi\mathbf{e}_h^*) \quad (\text{C.38})$$

Let us express the spatial frequencies measured by an array as two-element vectors of the form (ω_x, ω_y) . Let \mathbf{X} be the $d \times 2$ matrix containing these spatial frequencies. Note then that the phase-wrap error will manifest itself merely as an image shift if and only if this error is a (modulo- 2π) phase ramp, i.e. there exists a 2-element shift vector \mathbf{z} and an integer vector \mathbf{k} which together satisfy:

$$2\pi\mathbf{e}_\sigma - 2\pi\mathbf{X}\mathbf{z} = 2\pi\mathbf{k} \quad (\text{C.39})$$

Substituting from Equation (C.38) we obtain:

$$\mathbf{C}_{oc}^+(2\pi\mathbf{e}) - 2\pi\mathbf{X}\mathbf{z} = 2\pi\mathbf{k} \quad (\text{C.40})$$

Dividing through by 2π we obtain the equation: $\mathbf{C}_{oc}^+\mathbf{e}_h^* - \mathbf{X}\mathbf{z} = \mathbf{k}$. Note that each element of \mathbf{C}_{oc}^+ can be expressed as some rational number $\frac{p_i}{q_i}$. Similarly we first assume \mathbf{X} contains rational spatial frequencies with greatest common denominator q_x . Then we can multiply through by the least-common-multiple (LCM) of the $\{q_i\}$ and q_x to obtain a system of equations whose coefficients are guaranteed to be integer (i.e., we have a linear Diophantine

system). Let this LCM be denoted as l . Then we have, after rearranging terms,

$$l\mathbf{X}\mathbf{z} = l(\mathbf{C}_{oc}^+ \mathbf{e}_h^* - \mathbf{k}) \quad (\text{C.41})$$

We now wish to determine conditions under which there exist vectors \mathbf{k} and \mathbf{z} satisfying this overdetermined Diophantine system. Applying the Smith Normal Form decomposition (c.f. Theorem 2.2) to the matrix $l\mathbf{X}$ this time, and noting that $\text{rank}(\mathbf{X}) = 2$, we have:

$$\mathbf{D}_X = \mathbf{U}_X(l\mathbf{X})\mathbf{V}_X \quad (\text{C.42})$$

where \mathbf{U}_X and \mathbf{V}_X are unimodular matrices of size $d \times d$ and 2×2 , respectively, and \mathbf{D}_X is a rectangular diagonal matrix whose entries are zero below row 2.

If we left-multiply Equation (C.41) by \mathbf{U}_X on both sides, we obtain:

$$l\mathbf{U}_X\mathbf{X}\mathbf{z} = l\mathbf{U}_X(\mathbf{C}_{oc}^+ \mathbf{e}_h^* - \mathbf{k}) \quad (\text{C.43})$$

Using Equation (C.42) and the fact that \mathbf{V}_X is a unimodular (and hence invertible) matrix, we can then write:

$$\mathbf{D}_X\mathbf{V}_X^{-1}\mathbf{z} = l(\mathbf{U}_X\mathbf{C}_{oc}^+ \mathbf{e}_h^* - \mathbf{U}_X\mathbf{k}) \quad (\text{C.44})$$

We are now in position to prove the main result of this section, which is preceded by the following Lemma:

Lemma C.2: Given wrap-invariance, the (column) vector $\mathbf{U}_X\mathbf{C}_{oc}^+ \mathbf{e}_h^* = \mathbf{U}_X\mathbf{e}_\sigma$ has integer entries below row 2.

Proof:

Given that the elementary divisors of \mathbf{C}_{oc} are all 1, we know there exists an integer vector \mathbf{k}_0 such that:

$$\mathbf{e}_h^* = \mathbf{C}_{oc}\mathbf{k}_0 \quad (\text{C.45})$$

Substituting Equations (C.45) and the pseudo-inverse definition in (4.63) into Equation

(C.38), we obtain:

$$\mathbf{e}_\sigma = \mathbf{V}_\sigma \boldsymbol{\Sigma}_\sigma^+ \mathbf{U}_\sigma^T \mathbf{U}_\sigma \boldsymbol{\Sigma}_\sigma \mathbf{V}_\sigma^T \mathbf{k}_0 \quad (\text{C.46})$$

Noting that \mathbf{U}_σ is orthogonal, this equation can be simplified to

$$\mathbf{e}_\sigma = (\mathbf{V}_\sigma - \mathbf{N}) \mathbf{V}_\sigma^T \mathbf{k}_0 \quad (\text{C.47})$$

where \mathbf{N} is a matrix of the same size as \mathbf{V}_σ , which is zero except for the last two columns. These last two columns are identical to those of \mathbf{V}_σ , and hence by Lemma 4.1 comprise an orthogonal basis for the nullspace of \mathbf{C}_{oc} . Noting the orthogonality of \mathbf{V}_σ , this can be further simplified to:

$$\mathbf{e}_\sigma = \mathbf{k}_0 - \mathbf{N} \mathbf{V}_\sigma^T \mathbf{k}_0 \quad (\text{C.48})$$

Recall from Proposition 4.4.2 that the canonical basis for the nullspace of \mathbf{C}_{oc} is given by vectors $\Delta \mathbf{r}_x$ and $\Delta \mathbf{r}_y$, which denote the vectors containing the x - and y -coordinates of the baselines in an array, respectively. Note that each of the two non-zero column vectors $\{\mathbf{v}_k\}, k \in 1, 2$ in \mathbf{N} can be expressed (via a simple change of basis) as a linear combinations of the elements of the canonical basis as.

$$\mathbf{v}_k = a_1 \mathbf{w}_1 + a_2 \mathbf{w}_2 \quad (\text{C.49})$$

Hence we have:

$$\mathbf{U}_\chi \mathbf{e}_\sigma = \mathbf{U}_\chi \mathbf{k}_0 - \mathbf{U}_\chi \mathbf{N} \mathbf{V}_\sigma^T \mathbf{k}_0 \quad (\text{C.50})$$

Since \mathbf{U}_χ is an integer matrix, the first term is clearly integral. Let us then examine the second term, and in particular, the product $\mathbf{U}_\chi \mathbf{N}$. By substitution from Equation (C.49), we see that for any of the two non-zero columns \mathbf{v}_k of \mathbf{N} , we have:

$$\mathbf{U}_\chi \mathbf{v}_k = a_1 \mathbf{U}_\chi \mathbf{w}_1 + a_2 \mathbf{U}_\chi \mathbf{w}_2 \quad (\text{C.51})$$

Now note from Theorem 3.5.2 that \mathbf{U}_X is a matrix which annihilates all the spatial frequencies in the matrix \mathbf{X} below row 2. But these spatial frequencies are in fact the contents of the two columns \mathbf{w}_1 and \mathbf{w}_2 (up to a uniform scaling factor). Therefore the column vector in Equation (C.51) is zero below row 2. This means in turn that the second term in Equation (C.50) is zero below row 2, and hence that $\mathbf{U}_X \mathbf{e}_\sigma$ is integral below row 2 (since the first term is integral). \square

To utilize Lemma C.2, we first re-arrange the Equation (C.44) above so that it reads:

$$\frac{1}{l} \mathbf{D}_X \mathbf{V}_X^{-1} \mathbf{z} - \mathbf{U}_X \mathbf{C}_{oc}^+ \mathbf{e}_h^* = -\mathbf{U}_X \mathbf{k} \quad (\text{C.52})$$

Let $\mathbf{v} = \frac{1}{l} \mathbf{D}_X \mathbf{V}_X^{-1} \mathbf{z} - \mathbf{U}_X \mathbf{C}_{oc}^+ \mathbf{e}_h^*$. Note that since \mathbf{D}_X is zero below row 2, the entries of \mathbf{v} below row 2 will be equal to those of $(-\mathbf{U}_X \mathbf{C}_{oc}^+ \mathbf{e}_h^*)$, which are integers by Lemma C.2. Now consider the first and second entries of \mathbf{v} . Let \mathbf{f} be the vector containing the fractional parts of the first two elements of vector $\mathbf{U}_X \mathbf{C}_{oc}^+ \mathbf{e}_h^*$, and let \mathbf{A} be the invertible matrix consisting of the first two rows of $\frac{1}{l} \mathbf{D}_X \mathbf{V}_X^{-1}$. Without loss of generality, choose $\mathbf{z}^* = \mathbf{A}^{-1} \mathbf{f}$ so that the fractional part \mathbf{f} is annihilated, leaving only integer elements in the first two entries of \mathbf{v} . Hence we now have:

$$\mathbf{v} = -\mathbf{U}_X \mathbf{k} \quad (\text{C.53})$$

with \mathbf{v} ensured to contain only integer elements. Since \mathbf{U}_X is unimodular, the vector $\mathbf{k}^* = -\mathbf{U}_X^{-1} \mathbf{v}$ will be integral. We have thus found a pair $(\mathbf{z}^*, \mathbf{k}^*)$ with integer \mathbf{k}^* which satisfies the Equation (C.44). Since Equation (C.44) is related to Equation (C.41) via a unimodular (and hence invertible) mapping \mathbf{U}_X , invariance is hence proven. \square

C.4 Minimum Cycle Basis Proofs

Lemma C.4.1. (*Liebchen and Rizzi, 2005*): Let \mathcal{H} be the Horton family of a graph in which there is a unique minimum path between each pair of nodes. Let C be a cycle of G which is not in \mathcal{H} . Then there exists a minimum path $P_{u,v}$ between nodes in u and v of C which is internally disjoint from C .

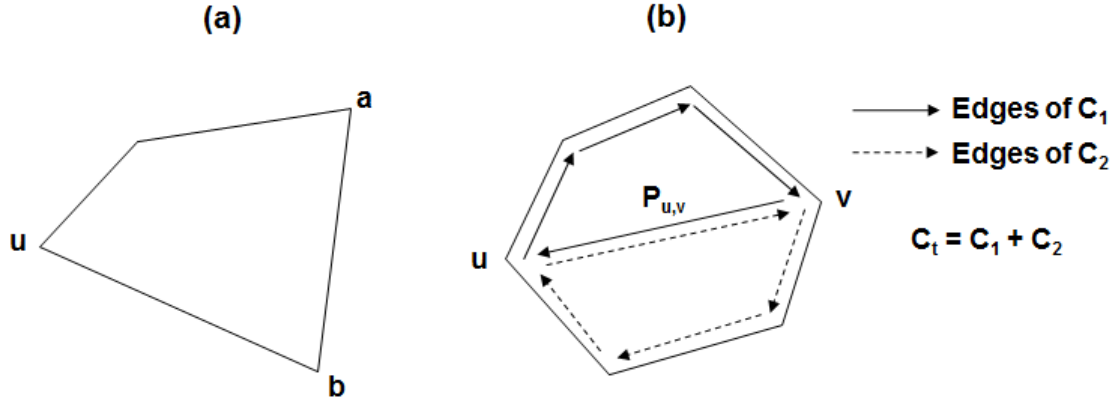


Figure C.1: Illustrating Lemma C.4.1 (a) and Lemma C.4.2 (b)

Proof: Choose an edge ab in C connecting nodes a and b (c.f. Figure C.1a), and another node u also in C . Since C is not a member of \mathcal{H} , either the path P_{ua} or P_{ub} in C must be sub-optimal (as otherwise $C \setminus ab$ would be a shortest-path tree rooted at vertex u). \square

Lemma C.4.2. *Liebchen and Rizzi (2005): All oriented cycles in a minimum cycle basis of a directed graph D are in the Horton family.*

Proof: Suppose we have found such a minimum cycle basis $\mathcal{B} = C_1, \dots, C_t, C_\mu$ and it was obtained by applying the greedy algorithm to the complete set of cycles of D . Furthermore without loss of generality assume that C_t is the first cycle which is not a member of the Horton set.

We know there exist two nodes u and v in C_t such that the shortest path $P_{u,v}$ between them is internally disjoint from C_t . Let C_1 and C_2 denote the two cycles in $C_t \cup P_{u,v}$ distinct from C_t and with opposite orientations on $P_{u,v}$, as shown in Figure C.1b. Clearly $w(C_1) < w(C_t)$ and $w(C_2) < w(C_t)$, and therefore both C_1 and C_2 can be expressed as linear combinations of the cycles in $\{C_1, \dots, C_{t-1}\}$. Note that $C_t = C_1 + C_2$, which implies then that C_t can be expressed as linear combinations of the cycles in $\{C_1, \dots, C_{t-1}\}$. But this is impossible since the greedy algorithm chose C_t as linearly-independent of these cycles, and hence we have the necessary contradiction. \square